AD-A212 269 MASTER COPY - FOR REPRODUCTION PURPOSES (2)

REPORT DOCUMENTATION PAGE

DTIC FILE COPY

DTIC ELECTE SEP 14 1989 S D B

1a. REPORT SECURITY CLASSIFICATION: Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT: Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S): ARO 22908.5-PH
6a. NAME OF PERFORMING ORGANIZATION: Stanford University; 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION: U. S. Army Research Office
6c. ADDRESS (City, State, and ZIP Code): Information Systems Laboratory, Stanford, CA 94305	7b. ADDRESS (City, State, and ZIP Code): P. O. Box 12211, Research Triangle Park, NC 27709-2211
8a. NAME OF FUNDING/SPONSORING ORGANIZATION: U. S. Army Research Office; 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER: DAAG29-85-K-0211
8c. ADDRESS (City, State, and ZIP Code): P. O. Box 12211, Research Triangle Park, NC 27709-2211	10. SOURCE OF FUNDING NUMBERS: PROGRAM ELEMENT NO. / PROJECT NO. / TASK NO. / WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification): Optical Clock Distribution to VLSI Chips

12. PERSONAL AUTHOR(S): Joseph W. Goodman

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final	FROM 8/1/85 TO 6/30/87	July, 1989	143

16. SUPPLEMENTARY NOTATION: The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

17. COSATI CODES		
FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number): Optical Clocks, VLSI Chips, CMSO Chips, Digital Logic Electronic Clocks

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Investigations of the use of optical for distributing the clock to a CMOS chip are reported. Two chips were designed, the first incorporating a set of integrated detectors followed by transimpedance amplifiers. The clock was distributed as an optical square wave. The detected clock signals were then amplified, distributed on polysilicon, and applied to digital logic. The chips were fabricated at the MOSIS facility. Testing of this family of chips

(over)

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS	Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE (Include Area Code) / 22c. OFFICE SYMBOL

yielded a maximum clock rate of about 20 MHz. A second chip was designed in which the optical clock was used only to synchronize a series of free-running electronic clocks distributed about the chip. Again the fabrication was performed at MOSIS. While theory and simulations predicted a maximum clock frequency of 100 MHz, it proved impossible to actually lock the clocks, due to nonuniformities in the fabrication parameters across the chip. Future work should aim at modifying the design for greater tolerance to fabrication nonuniformities.

ARO 22908.5-PH

FINAL REPORT

Optical Clock Distribution to VLSI Chips

Contract No. DAAG29-85-K-0211

To:

Army Research Office
Durham, NC
Attn: Dr. Bobby Guenther

From:

Information Systems Laboratory
Stanford University
Stanford, CA 94305

Joseph W. Goodman, Principal Investigator

July, 1989

ABSTRACT

Investigations of the use of optical for distributing the clock to a CMOS chip are reported. Two chips were designed, the first incorporating a set of integrated detectors followed by transimpedance amplifiers. The clock was distributed as an optical square wave. The detected clock signals were then amplified, distributed on polysilicon, and applied to digital logic. The chips were fabricated at the MOSIS facility. Testing of this family of chips yielded a maximum clock rate of about 20 MHz. A second chip was designed in which the optical clock was used only to synchronize a series of free-running electronic clocks distributed about the chip. Again the fabrication was performed at MOSIS. While theory and simulations predicted a maximum clock frequency of 100 MHz, it proved impossible to actually lock the clocks, due to nonuniformities in the fabrication parameters across the chip. Future work should aim at modifying the design for greater tolerance to fabrication nonuniformities.

I. Introduction and Overview

This document represents the final report on Army Research Office Contract No. DAAG29-85-K-0211, summarizing the work accomplished from August 1, 1985 through June 30, 1989. The report is organized into five sections, the first being this introduction and overview. Section II is a report on the most recent work on this contract, covering the period September 1988 through the contract end. The work during this period was devoted to testing the most recent chips, obtained from the DARPA MOSIS facility. This section of the report was written by Karin Sperley and A. Tabibian, both Stanford students. Section III summarizes certain administrative aspects of the grant.

Appendix A is a reproduction, in entirety, of the thesis of Bradley Clymer, who lead the work on the design of the first optically clocked chip built under this program. Bradley Clymer received his doctorate in Electrical Engineering, and is now an Assistant Professor at the University of Ohio. Appendix B is a reproduction, again in entirety, of the thesis of Richard Welch, who designed the second chip produced under this contract. Mr. Welch received the degree Engineer and is currently a student in the MBA program at the University of Illinois.

Thus to read the work in chronology, start with Appendix A, then read Appendix B, then finish with Section II.

II. Testing of The Optical Clock Distribution II Chip

A. Introduction

This section describes the testing of Mr. Welch's Optical Clock Distribution II (OCD2) chip. The testing was broken into two parts: testing without the optical clock signal present, and testing with this signal present.

B. Testing Without The Optical Clock Signal

The pre-testing of the chips consisted of checking the signals at the output pins of the powered chip to determine if the electrical circuitry was functioning properly. These outputs include those from the skew detector, frequency divider, and shift register circuitry. The pre-testing is described in section 3.2 of appendix B.

The results of the pre-testing, which was performed by Mr. Welch, indicated that nearly all of the skew pins and error pins from the shift registers on the chips had the correct signals present. However, most of the outputs from the frequency divider circuitry were incorrect. The problem was twofold. No signal was present in many cases, and when the signal was present the frequency was nearly half that predicted by the design. These problems are described in more detail below.

Each of the nine optical receiver/clock amplifier (ORCA) units on a chip generate a clock signal with a voltage controlled oscillator. From this signal two clock phases and their complements are generated. The frequency of the four phases of the selected ORCA are divided by 64 by the frequency divider circuitry and can be viewed on the four f_{out} pins. The problem was that the third and fourth phases of the clock were not present on the f_{out_2}

and f_{out3} pins of any chip, while the other two clock phases were sometimes present on f_{out0} and f_{out1}. The diagram below shows the physical location of the ORCAs on a chip. Each location is divided into four squares corresponding to the four clock phases generated from that ORCA's oscillator. The number in the box is the percentage of chips on which that particular phase of that particular ORCA was seen to be operational on an f_{out} pin.

Diagram of Operational Clock Phases and ORCAs

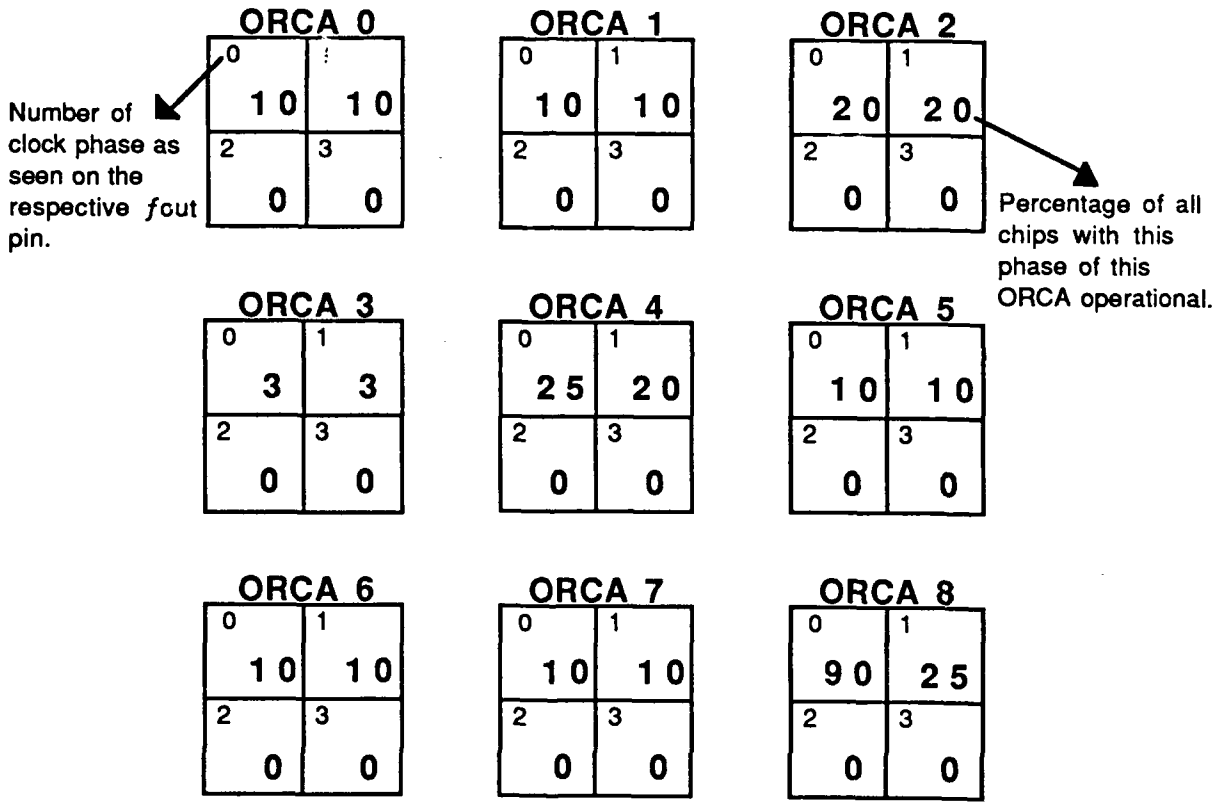

Figure 1.

Furthermore, the frequencies seen on the f_{out} pins correspond to oscillator frequencies of 50-65 MHZ, roughly one-half to two-thirds of the 100 MHz value predicted by SPICE models.

An interesting characteristic of the chips' operation was observed after the chips had warmed up. The signal output observed on an f_{out} pin became noisy in amplitude and phase. The signal became clear again if the supply voltage was lowered, but would again become noisy at the lowered voltage after a time. Also, lowering the voltage lowered the oscillator frequency.

The results of the pre-test indicated that Mr. Welch's design did not sufficiently account for fabrication processing variations since some of the circuits functioned and other identically designed circuits did not. In addition parts of the design do not work at all on any chip. It also seems likely that the SPICE models that were used may not be accurate at predicting the fabricated device performance since the designed operating frequency was nearly twice that of the actual device.

A simple test to check the correct operation and measure the transfer functions of the voltage controlled oscillators (VCOs) was performed in addition to the pre-tests. From the circuit diagram in Figure 2.11 in appendix B, it can be seen that in the absence of light input the output of the NAND gate (LF_Input) will always be high, i.e. equal to the supply voltage VCC. Referring to the circuit diagram of the low-pass/lag-lead filter in Figure 2.4, it can be seen that a DC LF_Input will result in a CNTR voltage equal to LF_Input. CNTR is the controlling voltage of the VCO, and therefore the VCO frequency may be tuned by adjusting VCC. This was done on ORCA2 of chip #4 and the resulting data appears in the graph below.

Note that as expected, the VCO responds linearly to changes in CNTR. The VCO ceased to operate at voltages below about 2 volts. Also note that the VCO can indeed achieve operating frequencies of 100 MHz and above, but this requires high VCC values at which the chips begins to overheat.

C. Testing With The Optical Clock Signal

The optical clock signal was generated with a 4mW laser diode at 750 nm wavelength. The laser was internally biased in its linear region and was directly modulated with a signal generator. The laser beam was delivered to the detector as shown in figure 3.

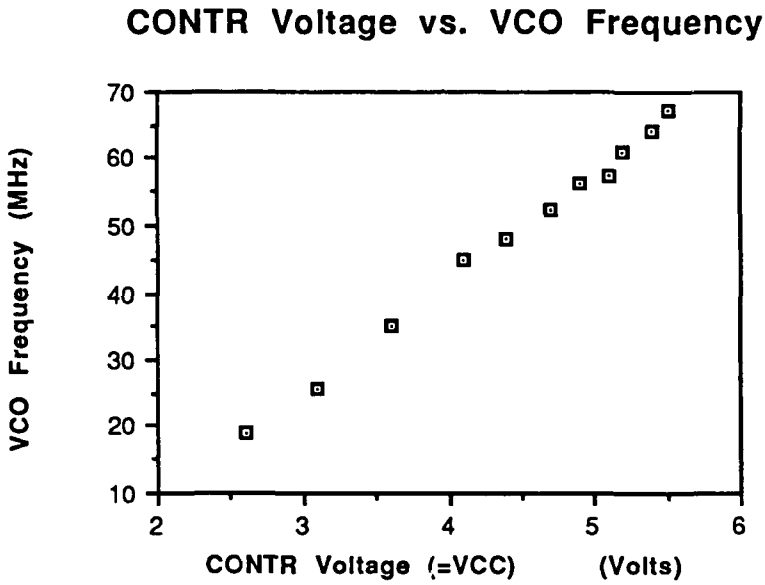

Figure 2.

Efforts to make the chip respond to the optical clock signal were fruitless. The photodetectors were exposed to laser light modulated at a few MHz below the natural operating frequency of the ORCA oscillators (50-65 MHz) and the frequency was then swept up in increments of 100KHz. The f_{out} pins were observed for signs of capture and lock of the Phase Locked Loop by direct observation of the oscilloscope trace and by externally triggering the scope with the modulation signal and watching for triggering of the f_{out} signal. However no change in output frequency or triggering were detected. The ORCAs' oscillators did not lock onto the modulating frequency of the laser. Different ORCAs on the various chips were tried with various laser power settings ; however, no PLL capture was detected. The only response to the laser observed on the f_{out} pins was a

small frequency shift corresponding to about a 1MHz drop in VCO operating frequency. This phenomenon is probably due to changes in device parameters caused by the heat of the incident laser. Since no ORCA locked onto the optical clock signal, the rest of the work consisted of testing the optical system for a possible source of this failure.

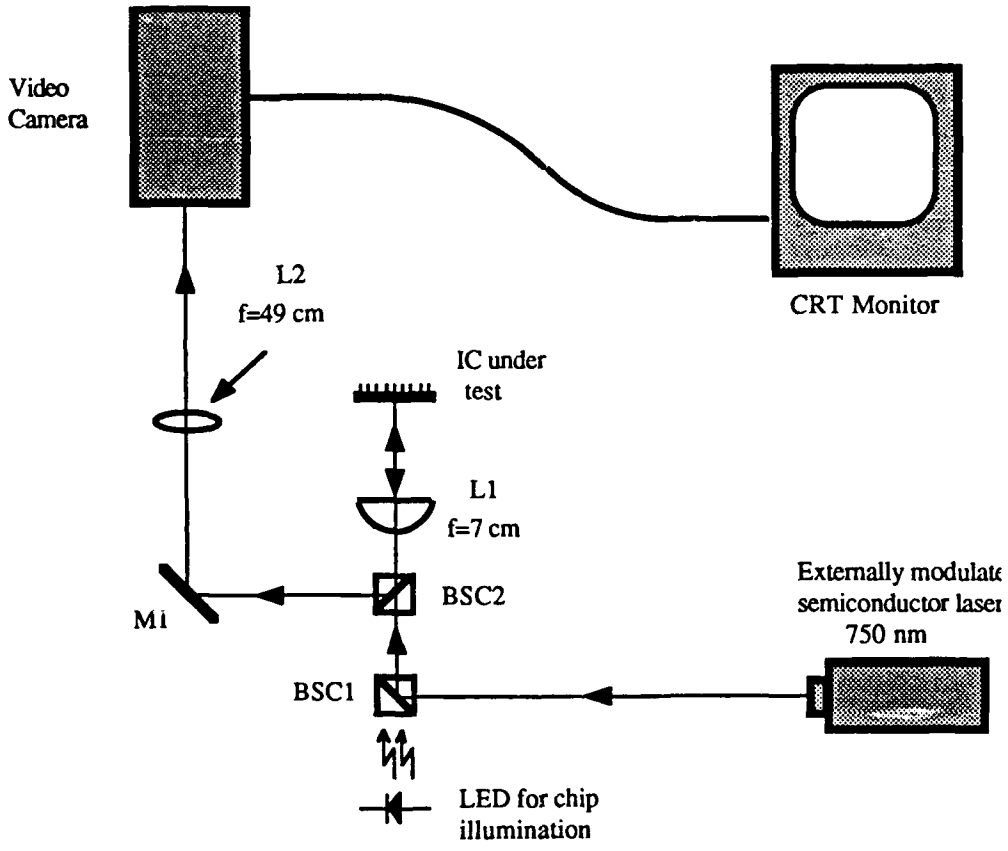

Figure 3

There are several possible causes for the failure of the ORCA's that deal with the external optical test equipment: The laser not modulating or not modulating at a frequency in the capture range of the phase-locked loop; the small, focused laser spot not hitting the detector; not enough or too much laser power on the detector; and the sinusoidal modulation of the laser is not sufficient to switch the digital circuitry (square-wave

modulation was assumed in the spice modeling of Mr. Welch). As many as possible of these causes for failure were isolated, tested, and eliminated as reasons for the chips' failure. The rest of this section describes this procedure.

The first concern was that the laser was not actually being modulated due to faulty circuitry in the laser or an impedance mismatch between the signal generator and the laser. It was found by utilizing a fast photodetector that impedance matching was required and that with this matching the laser was indeed modulated. To make sure that the modulation frequency was within the lock in range of the ORCA, the frequency was set as close to the natural frequency of the ORCA as possible. The external triggering technique was employed to detect the locking-on to the signal, but none was observed.

Alignment of the laser spot was a difficult procedure and very prone to error. This was due to the small size of the detector (20 microns square) and focused laser spot, and also because of the difficulty in viewing the 750 nm wavelength light focused onto a nearly specularly reflecting surface. To aid in alignment a microscope with a video camera and monitor was implemented. This was done by splitting off the light reflected from the chip's surface and focusing it with magnification onto a video camera. In order to view both the laser spot and the details of the chip's surface at the same time, it was necessary to both attenuate the laser light and to illuminate the chip's surface. In order to achieve the illumination without blocking the laser it was necessary to disassemble the delayed beam path (described by Mr. Welch in chapter 3 of appendix B) and insert a bright focusing LED next to the beamsplitter (see figure 3). This path removal was done to avoid laser power loss due to an additional beam splitter and because the delayed path was not needed until we could get the ORCAs to lock on the signal.

When viewed on the monitor the detectors could be seen and the laser spot could be approximately centered on the detector. It was difficult to estimate the spot size of the laser, however, due to blooming of the video camera. Even with careful alignment utilizing the monitor, the ORCA's still would not lock onto the laser signal.

Next the possibility of incorrect laser power was checked. Since no power meters were available, all that was done was to vary the average power of the laser and to watch for the ORCA to lock on. The ORCA was never observed to lock on at any laser power used.

The last possible cause of failure tested was the modulation signal. Initially a sinusoidal signal generator was used to modulate the laser because of availability. It became clear from reading Mr. Clymer's thesis (appendix A) that a square wave modulated clock signal was necessary. A pulse generator with sufficient bandwidth was obtained. However, the laser itself had a modulation bandwidth only extending to 500 MHz so that the laser limited the rise and fall times of the clock signal. The signal generator provided only enough power to modulate the laser about 60% which may or may not be sufficient (simulation on SPICE would be necessary). This pulse generator was used with visual alignment of the beam and various frequencies and powers were tried. Lock-on of the ORCA was never observed.

Since there seemed to be no point in continuing experimentation with laser modulation, testing with an unmodulated laser was attempted.

Sensitivity of the photodetector was tested by using different powers of laser light to the detector and measuring the chip's response via fluctuations in its supply current. Four different levels of incident light were used: 1) No incident light; 2) Room light; 3) Room

light + 1mW laser light; and 4) Room light + 2mW laser light. When the amount of incident light was changed a small but definite shift in the supply current was noticed.

However, moving the laser beam around in the vicinity of the photodetector showed that the maximum change in power supply current occurred when the laser beam appeared to be below and to the left of the photodetector. Therefore, these fluctuations in supply current may be due to the heating effects of the laser rather than changes in the state of the circuitry caused by current generated in the photodetector.

D. Theoretical Calculations And Simulations

Based on the expected component values of the low-pass/lag-lead filter a capture range of about 4MHz can be predicted. The procedure used for this calculation is taken from the Signetics 1985 "Linear Data Manual-Communications" pp. 4-236 to 4-256. Since Mr. Welch expected filter performance to be essentially constant under process variations, this number is expected to hold, at least theoretically. However, Mr. Welch's low-pass/lag-lead filter is a second order filter while the Signetics formulae assume a simple first-order filter. To a first approximation, however, the device values used in Welch's design are theoretically reasonable.

As mentioned before, the crucial problem with the SPICE simulations may be that they are not accurate enough at the 100 MHz frequencies used in this design. Validity of the diode models at 100 MHz frequencies is especially under question. Also, note that photodetectors are most often fabricated as p-i-n structures whereas Welch's MOSIS process did not allow the inclusion of an intrinsic layer thereby reducing the transport efficiency at the junction.

E. Conclusions And Suggestions For Future Work On Testing

Testing of Mr. Welch's chip did allow verification or near-verification of the operation of several of the designs components:

- By the mechanism described in the previous pages, it was possible to verify the correct operation of the Voltage Controlled Oscillator.

- It is safe to assume the correct operation of the frequency divider circuitry. Note that Figure 1 illustrates that none of the f_{out2} or f_{out3} pins show an output frequency. This cannot be due to flaws in the divide-by-64 circuitry design because the f_{out0} and f_{out1} use similar circuitry to divide the frequency of their outputs.

- The functionality of the low-pass/lag-lead filter for DC LF_Input values is certain.

- The design and operation of the phase generation circuitry has been verified despite the lack of output on the f_{out2} or f_{out3}. This can be said by considering that all the f_{out} frequencies are generated by the symmetric Phase Extraction Circuitry (Figure 2.15); therefore, the failure of two frequencies to appear on the pins is probably due to problems other than partially dysfunctional Phase Extraction Circuitry.

- As mentioned before the output of the Skew Detection and Error Detection pins was verified as correct under testing without optical input.

Despite these successes, perhaps the most instructive lesson of Mr. Welch's work has been to underscore the extra difficulties involved in testing a chip with an optical interface to the outside world.

A major cause of concern is that the current design affords no simple and effective method for the experimenter to determine the exact amount of light incident on the photodetector of each ORCA. The imaging techniques described in the Optical Testing section were used to direct the laser beam to the photodetector; however, the blooming due to the local saturation of the video camera's detector made it impossible to determine if the beam had indeed been entirely focused on the 20μm by 20μm photodetector. Hence, it was not possible to ascertain the amount of power incident on the exposed surface area of the photodetector. No external pin connections were included in the design that could have allowed this measurement. For future designs, it would be extremely prudent to provide pin connections or some other mechanism for determining the exact location of the beam on the chip. For instance, the main photodetector could be surrounded by other detectors with connections to outside the chip. By measuring the percentage of total light incident on each detector, it would then be possible to accurately calculate the position of the beam.

As far as Welch's chip is concerned, it would seem that the design philosophy of the chip is too optimistic. Before an entire operational chip can be built, the design of the individual sub-circuits must be understood and shown to be correct and reasonable. This is something that, depending on the sub-circuit, could not be done or was done with extreme difficulty on Welch's chip. Therefore, future designs must include far greater testing capabilities designed onto the chip. Obviously, these test structures must be as independent as possible from the other circuitry on the chip. Separate power supply connections, as well as a separate clock is recommended. The current design uses the four phases generated from a VCO to clock shift registers; failure to produce all four phases means that

much of the testing circuitry, such as the shift registers, will not operate. Many other testing schemes are possible, but in general a less aggressive, more modularized and testable approach is required.

The lack of proper equipment also proved to be a hindrance in the efficient testing of Welch's chip. In future research done at this or other universities, it might prove productive to arrange for testing to be done at the usually better equipped industrial laboratories.

III. Administrative Matters

During the period of this contract, the following papers and oral reports were presented on the work supported here:

B. Clymer and J.W. Goodman, "Optical clock distribution to silicon chips", *Proc. SPIE,* Vol. 625, pp. 134-142 (1986).

B. Clymer and J.W. Goodman, "Optical clock distribution to VLSI chips", *Optical Engineering,* Vol. 25, pp. 1103-1108 (1986).

B.D. Clymer and J.W. Goodman, "Timing uncertainty for receivers in optical clock distribution for VLSI", *Optical Engineering,* Vol 27, No. 11, pp. 944-954 (1988).

In addition, the work supported by this contract led directly to another project, supported by Digital Equipment Corporation, aimed at applying optical clock distribution to a higher level of the interconnect hierarchy. This work resulted in the following publication:

Ragai Khalil, Larry R. McAdams, and Joseph W. Goodman, "Optical clock distribution for high speed computers", Proc. SPIE, Vol. 991, pp. 32-41 (1988).

In addition, several oral papers on optical clock distribution were presented at annual meetings of the Optical Society of America during the duration of this contract.

During the duration of this contract, a number of different students worked on the clock distribution problem, including: Bradley Clymer, Richard Welch, Larry McAdams, Karin Sperley, and Ali Tabibian. One of these students (A. Tabibian) is an undergraduate, the others are or were graduate students.

APPENDIX A

OPTICAL CLOCK DISTRIBUTION FOR VLSI

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Bradley D. Clymer
September 1987

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

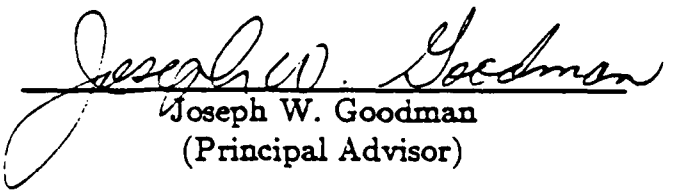

Joseph W. Goodman
(Principal Advisor)

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

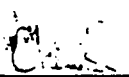

Lambertus Hesselink

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

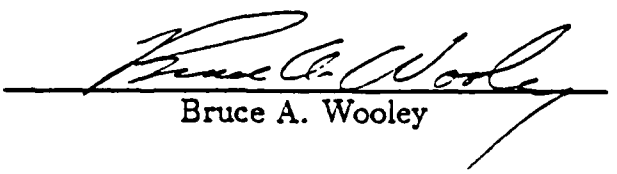

Bruce A. Wooley

Approved for the University Committee on Graduate Studies:

Dean of Graduate Studies & Research

In memory of

Juan Valenzuela

A friend who taught my spirit to move
and be moved

Abstract

Timing constraints for state-of-the-art very large scale integrated circuits (VLSI) in silicon are rapidly approaching communication limits available with layered two-dimensional metal and polysilicon wiring approaches. For such communication-limited systems, reliable clock distribution is a key concern. The range of finite differences in signal delays over clock wires of various lengths for large chips creates a timing skew that is significant when compared to the switching time of transistors in the circuit. The high bandwidth and three-dimensionality of imaging optical systems suggest that optical clock distribution systems have the potential to overcome the timing barriers presented by planar wiring. Clock signals can be holographically mapped to detector sites within small functional cells on a chip surface. Within each functional cell, the clock is distributed via surface wires with negligible delays. Since the difference in propagation time between optical paths is negligible when compared to typical electronic response times, the timing uncertainty for such a system is composed of two parts: the difference in response times (skew) between identical receiver copies implemented on a given chip, and the rms signal jitter at each receiver due to circuit noise. An overview of an optical distribution system is presented, along with system constraints and the design trade-offs these constraints represent. Two types of optical receivers for CMOS implementation are presented: a transimpedance receiver similar to those used in optical communication systems, and a phase-locked loop receiver which offers substantial improvement in performance and layout requirements over the first approach. Timing analysis for the two types of CMOS optical receivers is presented, showing that a 2-micron implementation of the transimpedance receiver

is limited to a maximum operating frequency of 50 MHz by fabrication-related response skews of up to 3 ns, while a 2-micron phase-locked loop receiver can operate in the 100 – 200 MHz range with steady-state response skews less than 55 ps and negligible signal jitter. A CMOS test chip to measure the performance of a 3-micron implementation of the transimpedance receiver has been fabricated. Laboratory measurements show that the maximum operating frequency of this receiver is approximately 15 MHz when 150 μW of incident optical power is applied, while response skew ranges from 5 – 20 ns. Observation of storage times for dynamic latch cells in proximity to photodiodes on the chip show that optical power and device separation specifications can be reasonably applied to guarantee minimum storage time.

Acknowledgements

The work that is presented in this thesis could not have been accomplished without the advice and support of many people during my study at Stanford, to whom I am very deeply indebted.

Perhaps the greatest experience I have had at Stanford is working under the supervision of Professor Joseph W. Goodman. His faith in my work and support through the many difficulties of doctoral study have been personal motivation for the completion of my research at Stanford. The example of his teaching, research and advisory style I shall carry as a role model for the rest of my career.

I gratefully acknowledge those many individuals with whom I have discussed my research problem and from whom I have benefitted greatly. Professor Bruce Wooley's suggestions led to many of the circuit designs and analysis presented in this thesis. Professor Ted Paige's discussions about optical interconnects and optical receivers have given many insights into the contributions and limitations of various approaches for optical signal distribution. The expertise of Dr. James Burr in VLSI design and the free use of the SETI group's layout and testing tools have made the realization of a test chip possible. The early discussion of photodetectors with Dr. James Sturm and amplifier circuits with Dr. Terry Walker provided the foundation for much of the work presented. Members of the research groups of Professor Goodman, Professor Hesselink and Professor Macovski have offered suggestions and advice for the optical layouts used in the laboratory procedures presented, especially Steve Collicott, Roy Matic, Robert McRuer, Richard Norgren and Jeff Wilde. John Acken, Anthony McCarthy, Timothy Schreyer and

J.T. Wu of the Integrated Circuits Laboratory have been especially helpful in providing device parameter values, packaging assistance and circuit development suggestions.

It has been my distinct pleasure to have been involved in the dance program during my period of study at Stanford. Through classes and performing groups in the dance division, I have learned creativity and expression which have become an integral part of my approach to other areas of my life as well. I would like to recognize the members of the Stanford Contemporary Dance Ensemble for their friendship, support and instruction, especially Jancy Limpert, Tony Kramer and my dear friend, Thane Kreiner.

My study at Stanford simply could not have been accomplished without the loving support of my wife, Susan. Not only did she comfort me when it seemed that all was lost, but she spent many evenings alone while I was writing my thesis and rehearsing for dance performances. There is not another woman alive that would have understood the problems that I have encountered in the last four years and also accept my methods for addressing them.

This work was support by a research grant from the Army Research Office.

Stanford University BRADLEY D. CLYMER
Spetember 1, 1987

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The large bandwidth, immunity to interference, and electrical isolation provided by optical communication techniques have led to recent interest in optical interconnections for computer systems. Optics has become an appealing alternative to wired interconnections on several levels of communication hierarchy within computing systems[1,2,3,4,5]. As timing performance for on-chip processing becomes communication limited, chip-wide clock signal distribution is rapidly becoming a problem in which an optical interconnection solution may be appropriate[1,4].

The miniaturization of integrated circuit elements by layout design scaling in very large scale integrated circuits (VLSI) has created a great deal of interest in the development of new methods for reducing the timing skew associated with transmitting signals via wires to remote locations on a chip[6,7]. As device sizes decrease and chip sizes increase with technological advances, the speed of instruction execution on a VLSI chip becomes limited by signal transmission delay rather than device switching delays[8,9,10]. Communication delay is especially critical in the distribution of the clock signal which is used to synchronize the operation of various devices on a chip. Parasitic transmission line capacitance and resistance over varying pathlengths for this widely distributed signal cause a skew in waveform arrival times at different locations on the chip.

For optical clock distribution, the clock signal is mapped at the speed of light from an off-chip laser diode via an optical element to photodetectors integrated on the surface of the silicon chip. In general, the optical approach could be three dimensional, using a hologram or lenslet array, or it could use integrated or fiber optics. The system presented here is three dimensional in that the space above the integrated circuit chip is used to route the optical signals rather than confining the light to fibers or waveguides in a planar or quasiplanar topology. In an optical system, clock skew can be reduced to essentially the variation in response times for the different receivers distributed over the chip surface.

1.1 Clock Distribution with Conventional Wire Paths

There is a design tradeoff inherently associated with wired distribution of signals on integrated circuit chips. This tradeoff on one hand involves the choice of conductor for the wire path, and on the other the availability of the medium chosen.

VLSI designers can route signals on three types of wires, each having a different resistivity and all having approximately the same capacitance per unit area. Aluminum wires have a resistivity that is two orders of magnitude less than polysilicon or diffusion paths, making aluminum a very desirable wiring medium, especially for long paths[11]. Chipwide distribution of the clock signal is generally required in VLSI designs, requiring long paths and suggesting metal as an appropriate medium.

Aluminum wires are also required for power and ground distribution because the resistivity of other media causes excessive voltage level degradation[12]. In addition, local connection between drain diffusions of p-channel and n-channel transistors in CMOS logic designs is made with metal wiring rather than polysilicon, to avoid forming junctions between the highly doped polysilicon wire and

one of the drain diffusions which is doped with the opposite type impurity. Since many VLSI fabrication technologies support only one or two levels of metal interconnection, the remaining metal available for global signal distribution can be severely limited. Distribution of timing signals in alumimum adds to the overhead requirements on the metal layers, and further reduces the space available for data paths.

Two problems with signal distribution by means of wires of finite resistance and capacitance become especially acute when the chipwide timing system is taken as an example. The clock signal is used to synchronize the operations of a very large number of devices on a VLSI chip. The number of devices that a clock distribution system must accomodate and the wide range of distances between devices create special manifestations of the general wired communication problems. The finite capacitance and resistivity of the wires represent a very large loading effect due to the extensive fanout of the clock distribution system. The wires to individual devices are ultimately interconnected in some parallel manner; therefore, the capacitances of the individual wires add to represent the load capacitance at the clock driver. The transition time for the output voltage of the clock driver is proportional to the ratio of the load capacitance to the driver capacitance. In this manner, a large capacitive load causes a broadening of signal pulses by increasing the transition time for the clock driver and slows the overall system operation.

One design method for reducing the effect of capacitive loading is the cascading of buffer stages with successively larger gate areas. This minimizes the total throughput delay by allowing each buffer to charge an optimum capacitive load. The optimum increase in gate area is by a factor of e; however, the scale by which the inverter size is increased is usually four or five in order to make efficient use of chip surface area[13].

Another clock distribution problem is associated with the finite resistance and capacitance of the signal paths. The circuit lines can be modeled as distributed RC paths, and as such, the waveform propagation is described by the diffusion

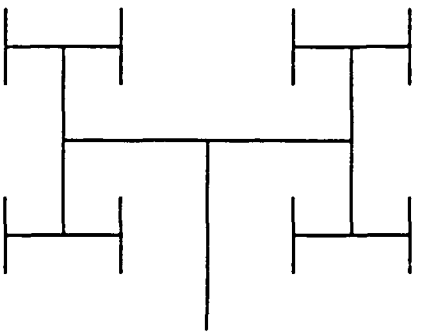

Figure 1: An H-tree clock distribution system.

equation[14]. Each length of wire has an associated delay that is a function of the resistance, capacitance and dimensions of the wire. Furthermore, as device sizes become smaller with layout design scaling, this characteristic communication delay increases quadratically for the same length of wire[9]. The large range of characteristic communication delays for the wires connecting the clock driver to the many clocked devices on a chip leads to a variation in the signal arrival time from one device to another. This is commonly referred to as clock skew, and it is a primary consideration in determining the system clock rate.

There have been several approaches suggested and implemented to reduce or eliminate the clock skew effect. One such approach, suggested by Anceau, involves distributing a lower frequency chipwide clock signal to several functional blocks, and internally synthesizing a high frequency clock to synchronize operations within each block[6]. A second approach involves forcing all wire lengths to be equal. One method for realizing this approach is an H-tree[7]. An example is shown in Figure 1. Line length variation effects are eliminated by the novel geometry of the layout. Remaining skew effects in such an H-tree structure are due to variations in the threshold voltages of the load transistors and characteristic capacitance and resistance between different paths[15]. Because the wire lengths are equal with this approach and chipwide distribution of the clock signal is required, H-tree distribution has the disadvantage of uniformly long wires, and therefore transitions are very slow due to a very large capacitive load for the clock

driver. In addition, the H-tree distribution requires a regular chip layout pattern, in opposition to the very popular modular design approach in which the layouts for individual functional cells are developed separately and incorporated in block form at the chip design level. A third clock distribution approach eliminates a chipwide synchronization signal by designing several functional blocks, each being self-timed. This approach allows fast execution of operations in the functional block, but at the expense of handshaking delays for communication and added control lines between functional blocks[16]. All of the approaches have the unfortunate characteristic of requiring massive use of metal wiring due to the lengths of communication lines necessary for signal coverage of an entire chip.

1.2 Optical Interconnection for Computing

Optics can provide several advantages for interconnection in computing systems[4]. One major advantage for optical distribution approaches is the freedom from mutual coupling between different signal paths. In Si-SiO_2 systems, there is a trade-off between distributed capacitance of an interconnection wire and mutual capacitive coupling of signals to adjacent wires. The capacitance between a wire and the ground plane can be reduced by increasing the dielectric thickness of the SiO_2, but in taking this action, the isolation between parallel wires is reduced, and increased mutual coupling results. Optical distribution systems do not exhibit a similar mutual coupling effect.

A second advantage of optics for computer interconnections is the ability of light beams to pass through each other without interacting. This is primarily due to the noninteraction of photons in linear media. This allows a more efficient use of layout space; it is valid for waveguide systems with angles of intersection greater than 10 degrees and virtually all free-space optical systems.

A third major advantage applicable to free-space optical interconnections is the potential to use rewritable optical materials to allow real-time reconfiguration of the interconnection pattern.

Research activities are currently in progress to incorporate optical approaches at nearly every level of computer interconnection hierarchy. Indeed, optical interconnections have already begun to appear for intercomputer local area networks[17, 18]. There is now activity to investigate optical interconnection between processors in a multiprocessor system[2]. Several programs have been applying optical solutions to interchip communication problems[19,20,21]. In addition, special optical distribution for signals with critical timing and large fan-out might be advantageous even though studies have shown that intrachip optical communication is not power efficient[3,22]. Clock distribution to several functional areas on a given VLSI chip represents this type of special signal distribution.

The implementations for optical interconnections that have been demonstrated or recommended are widely varying. Fiber optic systems have been demonstrated for local area networks between computers[17,18], and for lower levels of optical interconnection as well[19,21,23]. Other systems are based on planar waveguide optics[24] or three dimensional optics with holographic, grating or lenslet array elements [1,3,4,20,21,25]. While most applications reported in the literature have dealt with high speed data signals, several authors have suggested the use of optics for clock distribution. Hartman has proposed a waveguide distribution system at the board level, while Goodman, *et al.*[1,4], Fried[21], and Bergman, *et al.*[20] have suggested free-space optical systems for clock distribution at the chip level. Fried presents a phase-locked loop optical clock extraction circuit for CMOS based on a voltage controlled ring oscillator.

The optical clock distribution system that is presented here is a three dimensional design in which a clock signal is mapped from an off chip light source to several photodetector locations on an integrated chip surface. The beam mapping is performed by means of a holographic optical element or a lenslet array with negligible skew for the optical signal. The optical signal can be used in one of two manners. In one approach, the optical signal can simply be detected by means of a photodiode and amplified to a digital signal level, converted to a standard VLSI two-phase clock and distributed via short polysilicon wires to the individual

devices within each functional cell. In an alternative approach, the optical signal can be used to synchronize several local oscillators distributed on the VLSI chip. In this approach, the output of each local oscillator provides the input square wave for a two-phase clock synthesizer in each functional cell, and the two phases of the clock are again distributed via short polysilicon wires to devices in the functional cell. For either approach, clock signal skew is characterized by the variation in photoreceiver response times for identical receivers implemented on a given chip.

1.3 Organization

The next chapter is a detailed description of the proposed optical clock distribution system. An overview of the clock distribution system is presented first, followed by a description of the goals for the system design and implementation and a characterization of design trade-offs. The presentation includes two examples for division of the chip surface area into functional blocks along with suggested layout topologies within the cell for the optical receivers and the local polysilicon distribution system. The optical receiver circuit for each functional cell is expected to be one of those presented in Chapter 3 or Chapter 4.

A description and analysis of a standard transimpedance optical receiver is given in Chapter 3. This receiver is composed of analog amplifer stages and a digital phase divider circuit. Since a key source of clock skew for a multiple receiver chip is the difference in response times of identical receivers on a given chip, an extensive analysis is included which characterizes this variation as a function of fabrication-related parameter variation in the transistors on the chip. A second source of timing uncertainty is phase jitter. A noise analysis of the transimpedance amplifier is presented along with an example using typical device parameters for a 3-micron CMOS design.

Chapter 4 is a description and analysis of a phase-locked loop optical receiver. The elements of the phase-locked loop are described, including a voltage controlled ring oscillator and a novel biasing method for the photodiode to allow

phase detection without requiring a separate multiplier. A phase jitter analysis is also provided for the phase-locked loop receiver, followed by a comparison of the incident optical power levels required for the transimpedance and phase-locked loop receivers to achieve system timing uncertainty specifications.

The design of a CMOS chip to test receiver and photodiode performance parameters is presented in Chapter 5. Two tests involve measurement of the maximum operating frequency and response skew of a transimpedance amplifier receiver. A third test involves the measurement of leakage currents between the substrate under a photodiode and nearby dynamic memory devices. The specific circuit and measurement procedure are presented for each.

Finally, Chapter 6 contains a summary of contributions represented by the work that has been presented. In this chapter, suggestions are made for future investigation in the area of optical clock distribution systems.

Chapter 2

Optical Clock Distribution System Design

In this chapter, the proposed optical clock system is presented. A system overview is given first, followed by a description of the goals for the system design and implementation, and a characterization of design trade-offs. The presentation includes two examples for organization of the chip surface area into functional cells along with suggested layout topologies within the cell for the optical receivers and the local polysilicon distribution system. The optical receiver circuit for each functional cell is expected to be one of those presented in Chapter 3 or Chapter 4.

2.1 System Overview

An overview of the proposed optical clock distribution system is shown in Figure 2. The optical clock signal is generated by an off-chip light source, as shown at the top of the figure. The light source might be a directly modulated laser diode or LED, or an externally modulated continuous wave laser. The optical beam is mapped through a focusing element to various photodetector sites on the chip surface.

Figure 2: Optical clock distribution system overview.

The focusing element is expected to be a hologram or lenslet array, depending on the wavelength and coherence of the optical source chosen. The photodiodes in Figure 2 are shown in a regular pattern for simplicity. In practice, the location of the photodetectors is random, as determined by the modular design of the chip layout.

One of the advantages of this optical clock distribution system is that it allows for modular hierarchical layout design for the electronic circuits on the chip. In the system, the layout design of the chip is such that surface area is divided into small regions which we refer to as functional cells. The dimensions of the functional cell are determined by the maximum length of polysilicon wire for which signal propagation delay is acceptable for the given system operation requirements. Each such functional cell contains a photoreceiver for the detection of the optically distributed clock signal. Each photoreceiver includes the necessary electronics to create a digital two-phase clock from the single-phase optical input. The two-phase electronic clock is distributed via polysilicon wires to clocked devices within each functional cell. The use of polysilicon as a local distribution medium minimizes the requirements for aluminum in clock signal routing.

The optical clock distribution system that is presented represents a two level hierarchical distribution in which the top level of distribution is realized optically, and the bottom level is implemented via polysilicon wires. The characteristic communication delay for the optical path between the light source and the photodetectors on the surface of the chip is negligible when compared to the switching time of a typical electronic device on the chip. Since the length of the longest polysilicon wire is required to represent a negligible transmission line delay, the clock skew for the system is represented by the differences in response times for the various photodiodes and amplifiers that are distributed over the surface of the chip. If receiver designs can be found in which this type of skew is appreciably less than that for a wired distribution system for the same size chip, then the optical system represents an improvement over existing clock routing methods.

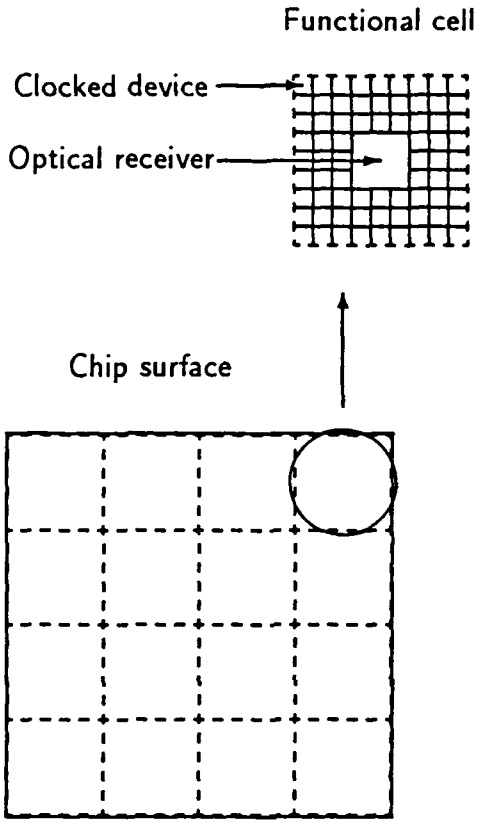

Figure 3: Chip layout for centered local clock distribution.

The layout organization of the functional cell can be adapted for the application expected. One organization involves placing the photoreceiver in the center of the functional cell, and distributing the clock signals in all directions to the surrounding devices. This approach is depicted in Figure 3. The inset shown in the upper right corner illustrates the organization of the functional cell in which the photoreceiver and clock driver are centered in the cell. The smaller squares surrounding the clock driver represent typical clocked devices, such as shift register stages.

A second organization might be more appropriate for dataflow layouts in which the clocked devices are arranged in lines for bit slice operations. For this type of application, the photoreceiver and clock driver can be designed to fit in a long

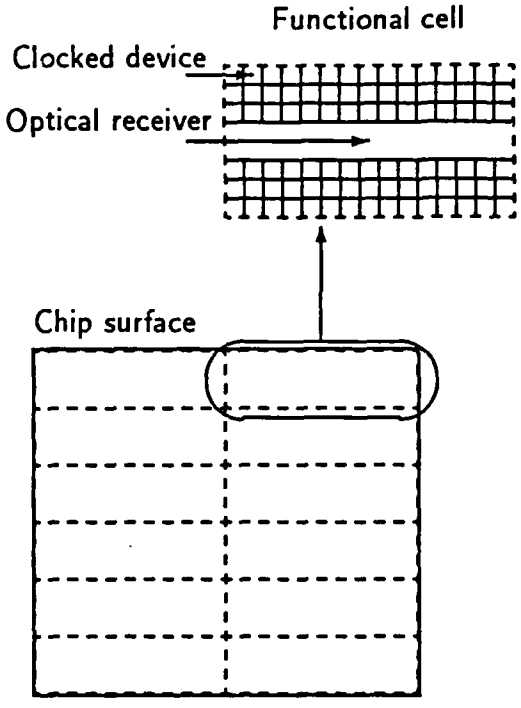

Figure 4: Chip layout for dataflow organization.

rectangular box, with the clock signals being distributed perpendicular to the length of the receiver. This design approach is shown in Figure 4. Again, the division of the chip into functional cells is shown in the lower left corner, while the organization of the functional cell is given in the inset at the upper right. The long box represents the photoreceiver and clock driver circuit, while the small squares represent the clocked elements in the dataflow path.

Two receiver designs have been developed which are compatible with optical detection and MOS technology. One design is a transimpedance amplifier; a detailed description of the design and operation of this receiver is presented in Chapter 3. The second receiver design is based on a phase-locked loop with a voltage controlled ring oscillator. A description of this circuit is given in Chapter

4.

2.2 Goals for system design and implementation

The design and implementation goals for the optical clock distribution system are primarily concerned with using existing technology whenever possible. In many cases there are conflicts that arise due to the difference in optical and electronic technologies and materials. The trade-offs that result from such conflicts are presented in the next section.

One of the principle design goals is to apply existing MOS technology for the photoreceiver circuit elements. This excludes the possibility of p-i-n photodiodes, because an intrinsic layer cannot be formed in standard MOS fabrication processes. Customized non-silicon photodiodes are also excluded, because packaging such devices would require hybrid mounting and bonding technology.

In standard MOS fabrication processes, the depth of the drain and source junctions is approximately 0.5 microns or less. The junction depth can affect the efficiency of a photodiode as well as the response time. Design trade-offs caused by fabrication contraints on photodiodes are discussed in further detail in Section 2.3.

A second goal for the optical receiver design is to use a photodiode bias within the range of supply voltages for the existing circuits on the chip. Since standard VLSI designs have supply voltages of 0 and 5 volts, the photodiode bias must lie in this range. This limits the width for the photodiode depletion layer, and ultimately the efficiency and response time of the device. Again, the consequences of this design choice are discussed in detail in the next section.

Another principle design goal is to use existing optical receiver designs if possible. Optical communications technology has developed good integrated receiver designs, some being implemented in MOS technology. An example of a MOS optical transimpedance receiver is given in [26]. The receiver described in Chapter 3 is based on this design. Other optical receivers based on phase-locked loops have

been reported in the literature[21,27]. The receiver described in Chapter 4 has been adapted from these designs.

A fourth goal for the optical clock distribution system is to use existing optical sources. A variety of sources are available from optical communication technology. The sources commonly used in optical fiber systems range from infrared laser diodes to visible LEDs. Many of the laser diode sources emit light at wavelength longer than the bandgap for silicon and are unusable for clock distribution; however, several inexpensive laser diodes emitting at wavelengths near 0.8 microns have been developed for optical disk storage systems. Another alternative could be using a visible wavelength laser and an acousto-optical modulator as the optical clock source. This approach seems applicable to large systems in which a single optical clock is distributed to several VLSI chips.

2.3 Design Trade-Offs

The principle system design trade-offs involve the physics of photodetection, the wavelength of the optical signal, the choice of optical source, and the choice of focusing element for the mapping of the optical signal to the photodetectors. A description of the relationship between these aspects of the optical clock distribution system is presented in this section.

In a reverse-biased p-n junction diode, an electric field is maintained across the depletion layer. Carriers generated in the depletion layer drift in the presence of this field across the diode junction to form a current, commonly called the drift current. Carriers generated outside the depletion layer diffuse from higher concentrations to lower concentrations. Since the depletion layer edge represents a lower concentration of minority carriers, minority carriers generated outside the depletion layer statistically tend to diffuse toward the depletion layer. If the carriers reach the edge of the depletion layer before recombining, then they are swept across the diode junction to form a current, commonly called the diffusion current.

The random nature of the diffusion process causes it to be a much slower mechanism than drift. For this reason, the drift current has a much faster response than the diffusion current. For fast photodetection, nearly all of the photons should be absorbed in the depletion layer, with the resulting photocurrent being primarily a drift current. For high sensitivity without a fast response, simply requiring that nearly all photons are absorbed with the characteristic recombination length of the depletion edge is sufficient [28].

The designer of a photodetector generally has the option of applying a large enough reverse bias to adjust the depletion layer width to match the photon absorption length for the wavelength of light that is detected. This is where a principle conflict between design elements of the goals for the optical clock distribution system is observed.

The photon absorption length in silicon for red light is around 2–3 microns. By applying a reverse bias of 5 volts, the depletion layer width can match this. If the prefered source technology is a laser diode, emitting at a near infrared wavelength of around 0.8 microns, then the photon absorption length in silicon is in the range of 10–25 microns. The depletion layer thickness can no longer be increased to match the absorption length because there is a maximum of 5 volts available for biasing the diode. The result is slower response times for the photodiode and poorer sensitivity. In addition, many of the photons that penetrate deep into the silicon substrate can diffuse to the depletion regions of nearby transistors rather than diffusing to the photodiode junction[29]. This not only represents a loss in efficiency, but also a crosstalk current. If the nearby transistor is an element in a dynamic storage device, the result could be a memory erasure with every clock cycle.

The problems of leakage currents, poor sensitivity and response time can be reduced by any of three methods. One choice involves using a laser diode source emitting light at a near infrared wavelength, and modifying the recombination characteristics of the silicon near the photodiode. This involves a non-standard fabrication process, and one of the design goals is compromised.

A second choice is to require a design rule specifying a minimum separation length between a photodiode and any other active device. This approach requires an increase in layout design constraints and overhead, and does not solve the poor response time and sensitivity problems.

The remaining choices involve using visible light for the optical signal. The designer can use a visible wavelength LED as the optical source, but light emitted by an LED is temporally and spatially incoherent, and a holographic focusing element requires coherent light for good resolution. A lenslet array can be used with the LED source, eliminating the requirement for coherent light, but the lenslet arrays that are commercially available are fabricated in regular patterns, and the random mapping ability of the focusing element is lost. If the photodiode locations are constrained to a regular pattern, the flexibility of the cellular layout design for the chip is compromised.

A final choice is to use a visible wavelength coherent source such as a He-Ne or argon ion laser. These lasers can be operated in a continuous wave mode, and the beam can be modulated externally by means of an acousto-optic modulator. In this approach, the packaging of the system becomes less attractive than with semiconductor sources, but perhaps the higher optical output power of the larger laser can be used to provide an optical clock signal to an entire system rather than a single chip to compensate for the added complexity.

The options of coherence, wavelength and packaging of the optical source for the clock distribution system presented allow the designer to choose the source that best fits the specifications of the VLSI system with which it will be implemented. The trade-offs that have been presented are characteristic of any optical clock distribution system in which the design goals of Section 2.2 apply, regardless of the optical receiver design. The specific receiver designs that are anticipated for use in an optical clock distribution system are presented in Chapter 3 and Chapter 4 and accomodate any of the optical source or focusing element choices that have been presented here.

Chapter 3

Transimpedance Receiver and Clock Driver

This chapter is a presentation of the design and operating parameters of a simple transimpedance optical receiver and clock driver. The receiver and two phase clock drivers circuits are both CMOS adaptations of NMOS designs that have appeared in the literature[26,30]. Descriptions of the receiver and clock driver circuits are given in Section 3.1.

The timing uncertainty in a distribution system composed of several receiver circuits consists of two parts. One type of uncertainty is static, *i.e.* it is fixed for a given chip and varies for different issues of the chip. This is the skew due to differences in receiver response times caused by fabrication-related variations in transistor channel dimensions and threshold voltages on a given chip. A second uncertainty is temporal and is characterized by phase variations due to noise in the receiver circuit. A computer simulation of the static timing skew for the transimpedance receiver and clock driver is presented in Section 3.2, and an analysis of the clock signal phase noise is given in Section 3.3.

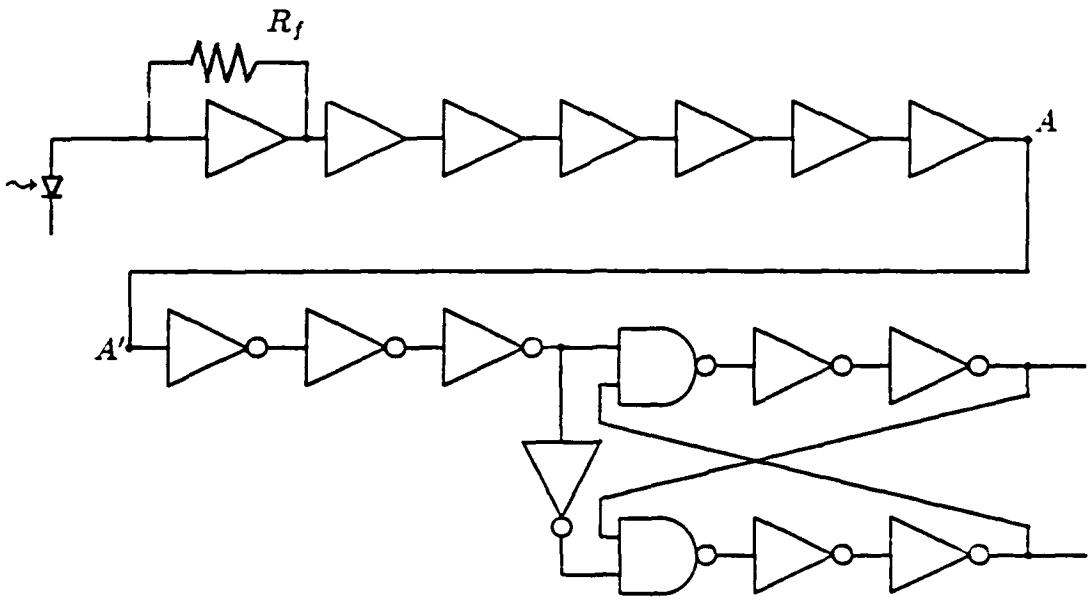

Figure 5: A transimpedance receiver and clock driver circuit.

3.1 Circuit Descriptions

A circuit diagram for a transimpedance optical receiver and clock driver is shown in Figure 5. The devices shown to the left of A in the figure represent a transimpedance receiver with analog inverting buffer stages to amplify the detected optical signal to voltage levels suitable for input to digital CMOS logic devices. The design and operation of this receiver are presented in Section 3.1.1. This design is very simple and is intended to demonstrate the performance nature of this type of receiver, rather than reflect an optimized circuit. The circuit elements to the right of A' represent a standard VLSI two-phase clock driver circuit. The details of the driver circuit are presented in Section 3.1.2 with a brief introduction to two-phase clock systems for VLSI.

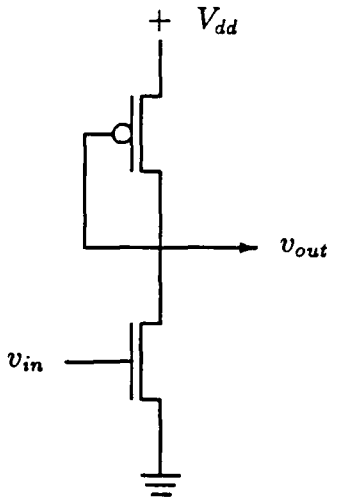

Figure 6: Analog CMOS inverting voltage amplifier stage.

3.1.1 Transimpedance Receiver

The transimpedance receiver that is shown in Figure 5 has been adapted for CMOS from a design published by Abidi[26]. The receiver circuit is composed of analog amplifier stages. Each stage contains an NMOS pull-down transistor and a passive PMOS pull-up resistor, as shown in Figure 6. The first stage has a resistive feedback path, as shown in Figure 5, which allows the photocurrent to modulate the gate voltage of the NMOS transistor of the first stage. The output of the first stage is modulated accordingly, and is the input voltage for the next inverter. The ratio of pull-up to pull-down channel length dimensions for all stages in the chain is constant and determined by the biasing requirements of the receiver.

In addition to providing a transimpedance path, the resistive feedback also serves to self-align the dark state operating points of the inverter stages in the chain. When no photocurrent is present, each inverter stage rests at a steady state operating point. Since there is no current flowing through the feedback resistor, the output voltage of the first stage is equal to the input voltage, and the inverter is held in equilibrium. Because the stages are designed to have similar dimensions

and the gain per stage is small, this equilibrium point is identical for all stages in the chain; therefore, when the input stage returns to equilibrium, all subsequent stages follow.

The photodetection and amplification can be described in terms of the modulation of the amplifier chain voltage between the dark state equilibrium and a lighted state level. When a photocurrent is present, the output voltage of the first stage is driven above the dark stage bias point, and the subsequent stages each amplify and invert this signal. When the light is removed, the output voltage of each stage returns to the dark state equilibrium point. Since the output of the amplifier chain is required to provide reliable TTL input voltage levels for the two-phase clock driver circuit, the dark state equilibrium voltage of the chain must be small enough to be reliably detected as a low logic level, and the lighted state output voltage of the last stage must be large enough to represent a high TTL input level. The choice of TTL voltage levels is a consequence of the design constraint to apply the optical clock distribution approach to standard CMOS VLSI systems.

The bias voltage also determines the gain for each inverter stage in that it determines the transconductance for each transistor. The open loop gain for a CMOS inverting amplifier stage is given as

$$A_v = -\frac{g_{m1}}{g_{m2}} = -\frac{\left[\frac{2I_D}{V_{GS1}-V_{t1}}\right]}{\left[\frac{2I_D}{V_{GS2}-V_{t2}}\right]} = -\frac{V_{GS2}-V_{t2}}{V_{GS1}-V_{t1}} \tag{1}$$

where g_{m1} and g_{m2} are the transconductances of the NMOS and PMOS transistors, respectively, I_D is the drain current, V_{GS1} and V_{GS2} are the gate-to-source voltages, and V_{t1} and V_{t2} are the threshold voltages[31].

A third amplifier design parameter that is affected by the bias of the transistors in the inverter stages is the maximum rate at which the output voltage of a given stage can make a transition from one voltage level to another. The is commonly called the slew rate, and often it determines the maximum operating frequency of a receiver. The relationship between the slew rate and the bias voltage for a

single inverter is given as

$$SR = (V_{GS} - V_t)\omega_1 \quad (2)$$

where V_{GS} and V_t are the gate and threshold voltages of the pull-down transistor and ω_1 is the unity gain frequency of the amplifier[32].

A comparison of Equation 1 and Equation 2 shows that a trade-off exists in determining the bias voltage. To maximize the open loop gain for each stage, the designer prefers a bias in which the gate voltage of the pulldown transistor is very near the threshold voltage. An amplifier stage with an optimum slew rate, however, should be designed so that the gate voltage is much higher than the threshold. For the circuit simulations presented in Section 3.2 and the circuits used in the test chip described in Chapter 5, a bias voltage has been chosen which permits a high slew rate at the expense of very low gain per stage. To some extent, the low gain per stage can be compensated by an increase in the number of stages in the amplifier chain. The gain per stage for these circuits is around 2, while the receiver slew rate is sufficient to amplify 67 MHz input signals to TTL voltage levels when transistor channel lengths are 2 microns.

3.1.2 Two-Phase Clock Driver

The output of the transimpedence receiver that is presented in the previous section is a single square wave signal. While a single-phase clock can be used to synchronize operations for computing, VLSI designers generally use two clock signals with mutually exclusive active phases. An introduction to timing and synchronization with two-phase clock signals and a simple circuit that can be used to synthesize the two phases from a single square wave are presented in this section. A more thorough description of timing for VLSI is found in Ref. [33].

In synchronous logic systems, the clock signal serves to maintain the order of occurance of events as well as defining the time frame in which combinatorial logic may be performed. During execution of a sequence of events, the clock signal controls an instruction counter which determines which operations are performed

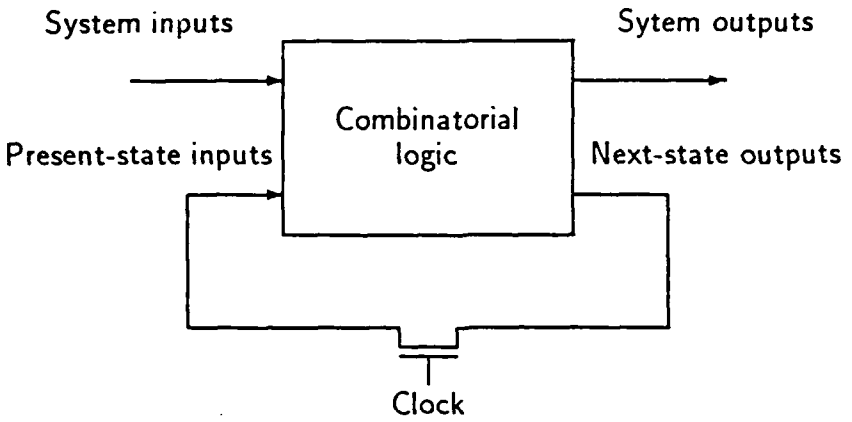

Figure 7: A finite state machine with a single-phase clock.

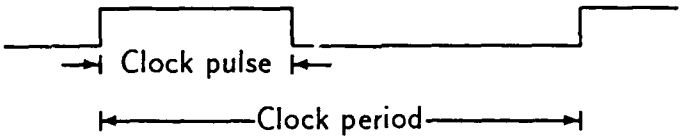

Figure 8: Timing for a single-phase clock.

on each clock pulse. In this manner, the execution speed of the processor is limited by the clock signal frequency. In addition, the width of each clock signal pulse defines the time frame in which a given set of instructions must be executed. The clock period must be long enough to guarantee that all of the required logic operations have settled to a steady state before the next pulse occurs.

The advantages of a two-phase clocking scheme are most easily demonstrated by first considering the operation of a finite state machine with a single-phase clock. Such a system is shown in Figure 7 and the corresponding clock signal is shown in Figure 8. When a positive clock signal is applied to the pass transistor, the next-state output voltage levels are transfered to the present-state inputs. The combinatorial logic operations begin as soon as the input voltages are presented,

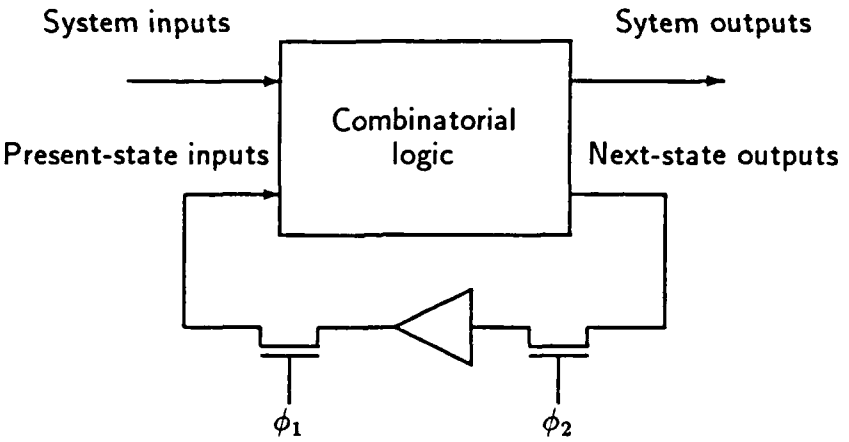

Figure 9: A finite state machine with a two-phase clock.

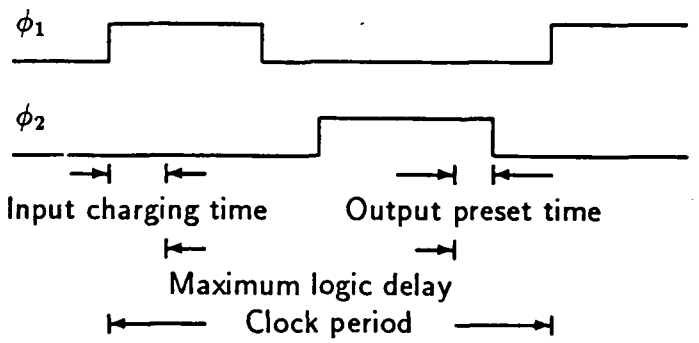

Figure 10: Timing for a two-phase clock.

and if the clock pulse is too wide, the output voltages due to the combinatorial logic execution can be reapplied to the inputs via the pass transistor. Since the next-state outputs become active inputs to the present state combinatorial logic devices before they are expected, errors in processing occur. For reliable operation, the combinatorial logic delay must be longer than the width of the clock pulse, yet shorter than the clock period.

The two-sided design constraint that is present in single-phase clock systems

can be reduced to a one-sided constraint by the use of two clock signals. This approach is commonly called two-phase clocking, and an example of a finite state system with a two-phase clock is shown in Figure 9. A timing diagram for the two-phases of the clock is shown in Figure 10.

An important relationship between the two clock phases, ϕ_1 and ϕ_2, shown in Figure 10, is that when one is positive the other must be zero. In the timing of this system, there are four time intervals: an interval in which only ϕ_1 is positive; an interval in which ϕ_1 returns to zero, allowing ϕ_2 to become positive; an interval in which only ϕ_2 is positive; and an interval in which ϕ_2 returns to zero, allowing ϕ_1 to become positive. During the positive interval of ϕ_1, previously stored data are transfered from intermediate buffers via pass transistors to present-state inputs. ϕ_1 must remain positive long enough to charge the input capacitance through the pass transistor. After charging the input capacitance, combinatorial logic is executed and output voltage levels reach a steady state. The width of the pulse for ϕ_1 can be much longer than the execution time of the combinatorial logic, because the feedback path is blocked by the zero value of ϕ_2. The transition of ϕ_1 to zero can occur at any time after the inputs have been charged, because the input voltage levels are maintained once the pass transistor is turned off. ϕ_2 may become positive at anytime after ϕ_1 has become zero, however it must remain positive long enough for the combinatorial logic to be completed and the output voltage levels for the next state to be transfered to the intermediate buffer by means of the second pass transitor. Since ϕ_1 and ϕ_2 are not allowed to be positive at the same time, there is no continuous feedback path; therefore, the only timing constraint is that the clock period must be sufficiently long to allow the combinatorial logic to be completed, along with a nominal delay for charging the present-state inputs and a preset time to insure stable next-state outputs.

It should be noted that for CMOS circuits, the pass devices are generally implemented as a PMOS and NMOS transistor connected in parallel, as shown in Figure 11. In this configuration, the control signal for the PMOS transistor is the complement of the control for the NMOS transistor. Four clock signals must be

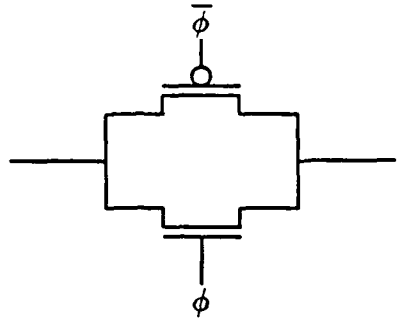

Figure 11: CMOS pass device.

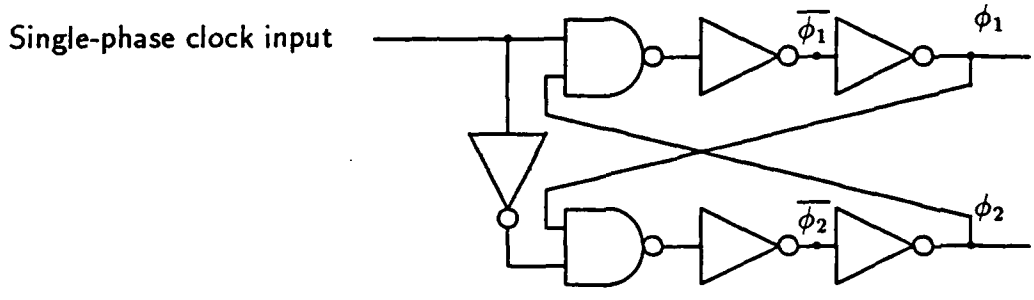

Figure 12: A two-phase clock synthesis circuit.

distributed for CMOS VLSI systems, *eg.* ϕ_1, ϕ_2 and the complements of ϕ_1 and ϕ_2.

Circuits which synthesize two-phase clock signals from a single square wave input are well understood. An example of such a circuit is shown in Figure 12. This clock driver is modeled after those discussed in Ref. [33]. The condition of mutually exclusive positive clock phases is achieved by using each clock phase as an input to a nand gate to generate the other clock phase. In this manner, the positive phase of one clock signal is dependent on the zero level of the other. The buffers following the nand gates in Figure 12 provide the drive capability to distribute the signal to the large capacitive load. The complementary clock phases are present at the inputs to the last buffer stage for each clock signal.

3.2 Static Receiver Response Skew Simulations

In a distribution system consisting of several receiver-driver circuits on a given chip, the signal skew can be represented as the difference in response times for the copies of these circuits on the chip. The designer can minimize the response skew effect by using identical layouts for the receiver-drivers; however, the actual transistor parameters can vary from one copy to another on a given chip due to the nature of the fabrication process. The fabrication-induced parameter variations cause differences in response times for identically drawn circuits and can be a major component in signal skew, especially for analog MOS circuits. The circuit simulation study that is presented in this section has been performed to determine the typical receiver-driver response skew resulting from reasonable variations in transistor channel dimensions and threshold voltages.

A SPICE simulation of the output signal for the simple optical clock detector and driver presented in Section 3.1 is shown in Figure 13. Here, the two clock phases are shown for an example with the nominal transistor threshold voltage and channel dimension values. The simulation reflects a layout design in which the nominal transistor channel dimension is 2 microns. The nominal width of each output buffer stage is 320 microns. The width of the input pull-down transistor for the transimpedance amplifier is 1380 microns and represents the largest device dimension in the circuit. As noted in Section 3.1, this circuit is intentionally designed to be simple and device dimensions have not been optimized with respect to trade-offs between output response and layout overhead. The experiment is intended to give an indication of operation of the general approach when photodetection and clock driver circuitry are confined to a reasonably small area (less than 0.003 cm^2).

For all simulations, the photodetector is modeled as a reverse biased p-n junction diode in parallel with an independent current source. The current source waveform is modeled as a 12.5 μA peak current square wave with linear transitions of 10 ns. The equivalent peak optical power incident on the photodetector

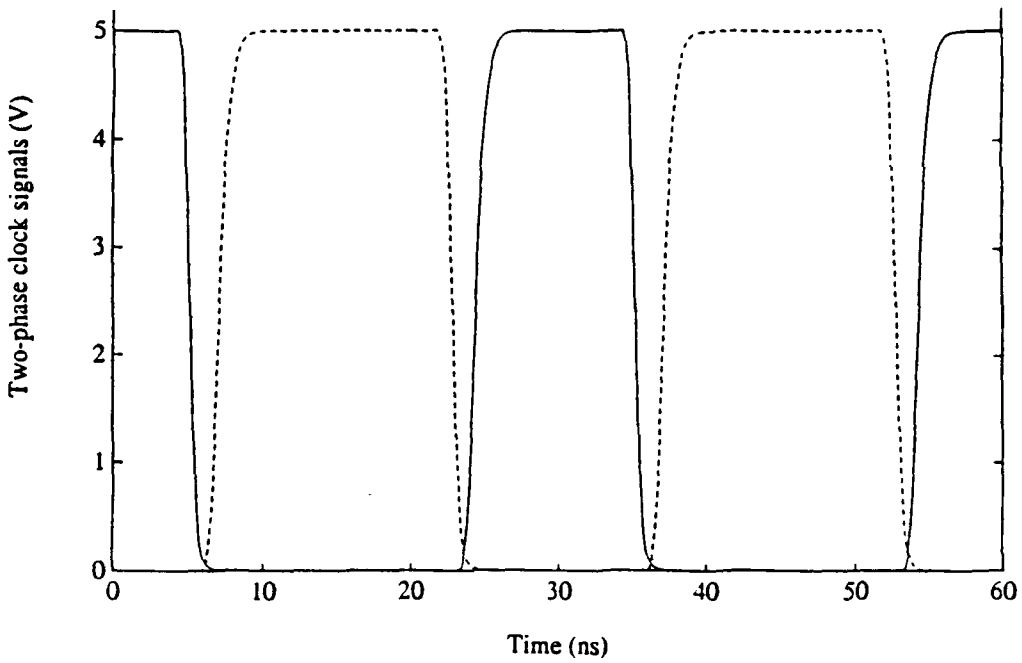

Figure 13: SPICE simulation of ϕ_1 (solid) and ϕ_2 (dashed) clock phases for trans-impedance receiver and two-phase clock synthesizer.

(for a square wave with a 50 per cent duty-cycle) which is expected to achieve such a current level is on the order of -10 dBm for visible wavelengths. The capacitive load is 2.5 pF, and represents the fanout for a typical functional cell described in Section 2.1.

Parameter variations between transistors which are closely spaced on a given chip tend to be highly correlated. For example, if an NMOS transistor has a channel dimension that is 0.1 μm longer than the ideal value, all NMOS transistors in that general area of the chip surface tend to have the same error. Parameter errors for devices separated by large distances on a given chip tend to be more independent. For the simulation studies presented here, all of the transistor parameters for similar devices in a single circuit simulation were varied the same

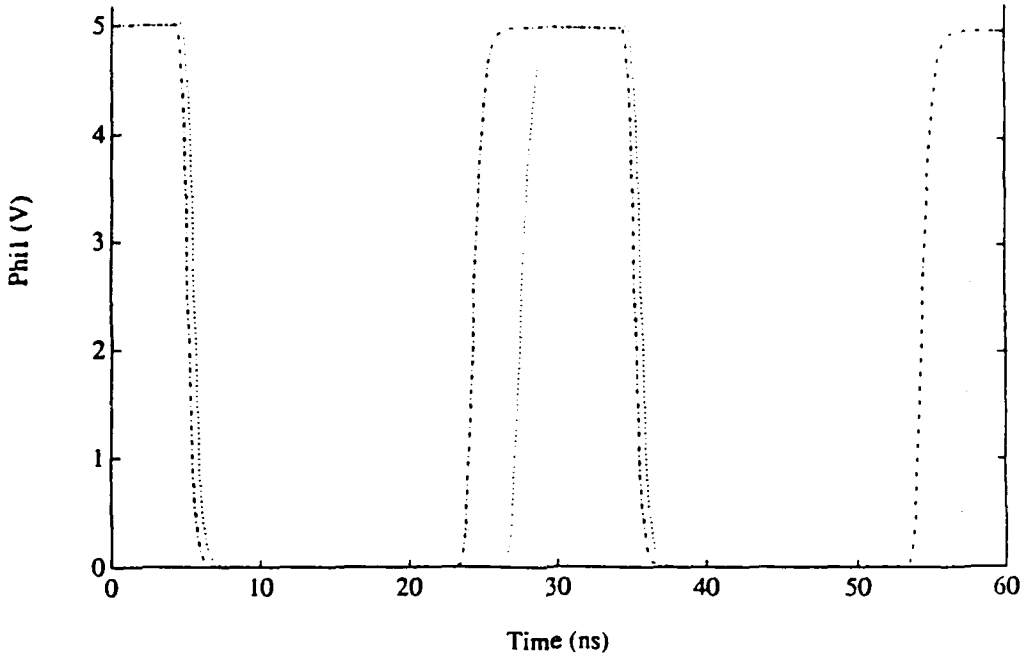

Figure 14: Worst case clock phase signal skew for transimpedance receiver and clock driver.

amount. Transistor dimension variations were assumed to be $\pm 0.1\mu$m and threshold voltage variations to be ± 1 mV, reflecting parameter fluctuations that might be expected on a given chip for a good 2-micron fabrication process. Channel width variations in these ranges caused no noticeable changes in the response of the receiver or clock driver output. The worst case output signal variations for changes in channel length and threshold voltages are shown in Figure 14.

In Figure 14, the signals represent the minimum and maximum throughtput delays for the output of the clock driver from point B in the circuit diagram of Figure 5. This is the ϕ_1 clock signal. The separation between curves represents the receiver skew that is expected due to fabrication-related differences in transistor parameters. The transition time for each signal is around 1 ns, and the maximum response skew is around 3 ns. These values are representative of a very simple

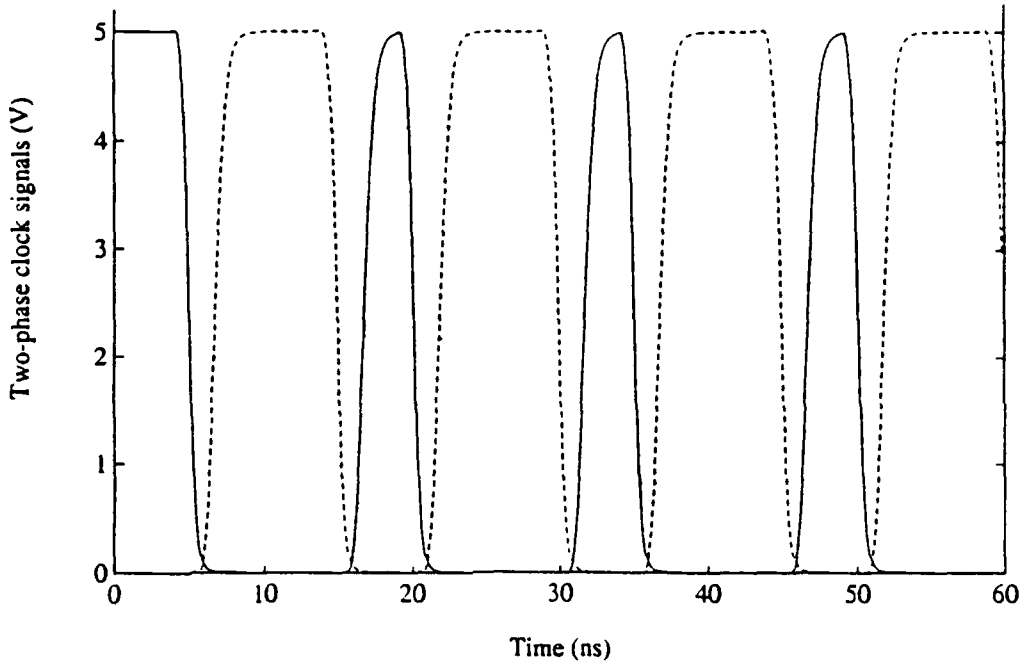

Figure 15: ϕ_1 (solid) and ϕ_2 (dashed) receiver-driver clock phases response for 67 MHz optical input signal.

photoreceiver and clock driver circuit which has not been optimized with respect to minimizing receiver skew.

In another series of simulations, the input optical clock frequency was varied to determine a maximum operating frequency for the circuit. This simulation assumes a circuit which is not limited by photodiode carrier diffusion delays, but rather provides an indication of the maximum frequency of the amplifier and clock driver electronics. The receiver and clock driver were found in the simulations to produce valid digital output levels for the two clock phases up to around 67 MHz, as shown in Figure 15; however, for reliable clock distribution, the expected receiver skew must be accounted for. This allows clock frequencies of up to 50 MHz, based on the simple circuit used in these simulations. A more sophisticated circuit design might easily achieve a higher clock rate.

3.3 Signal Jitter Analysis

The static receiver response skew that has been presented represents only one type of timing uncertainty that is inherent in a synchronous system with distributed clock receivers and drivers. For the static timing uncertainty, each receiver has a specific response delay, and the variation of delays between receivers results in a signal skew. A second type of timing uncertainty involves the temporal variation in the transitions of the clock signal at each receiver due to noise in the system. This type of timing uncertainty is commonly called phase jitter. This section is a presentation of the phase noise analysis of the transimpedance receiver described in Section 3.1.

Since phase noise is stochastic in nature, phase jitter is commonly expressed in mean square units of radians^2 or degrees^2. For purposes of comparing the stability of clock signals, it is sometimes more helpful to express the jitter in terms of time units rather than phase units. This is especially useful when comparing receivers that have different operating frequencies, as is done in Chapter 4. The analysis in this section will be first expressed in terms of phase units, then converted to time units for a typical frequency of operation that is expected.

In order to characterize the phase noise present at the output of the receiver, a signal-to-noise analysis technique for phase and frequency modulated systems from communication theory is employed[34]. The input signal is approximated by a sinusoidal waveform to simplify the analysis. The signal and noise currents at the receiver input are given by

$$i_{sig}(t) = \frac{I_{ph}}{2}(1 + \cos\omega_c t) \tag{3}$$

and

$$i_n(t) = i_{nc}(t)\cos\omega_c t - i_{ns}(t)\sin\omega_c t \tag{4}$$

where I_{ph} is the amplitude of the clock signal at the input of the detector, i_{nc} and i_{ns} are the quadrature representations of the bandpass noise, and ω_c is the frequency of the clock signal. From Equation 3 and Equation 4, the a-c portion

of the current at the input to the receiver can be represented as

$$\begin{aligned} i_{in}(t) &= \frac{I_{ph}}{2}\cos\omega_c t + i_{nc}(t)\cos\omega_c t - i_{ns}(t)\sin\omega_c t & (5)\\ &= [(I_{ph}/2) + i_{nc}(t)]\cos\omega_c t - i_{ns}(t)\sin\omega_c t & (6)\\ &= R(t)\cos[\omega_c t + \theta(t)] & (7)\end{aligned}$$

where $R(t)$ is the time-varying envelope and $\theta(t)$ is the phase variation due to noise, as determined by

$$\begin{aligned} R(t) &= \sqrt{[(I_{ph}/2) + i_{nc}(t)]^2 + [i_{ns}(t)]^2} & (8)\\ \theta(t) &= \arctan\frac{i_{ns}(t)}{(I_{ph}/2) + i_{nc}(t)}. & (9)\end{aligned}$$

In the standard analysis, as is the case for the receiver that has been presented, it is assumed that the received signal is eventually clipped; therefore, the time-varying envelope can be approximated by the time invariant amplitude of the signal after being limited, A_L. The output of the limiter is then given as

$$i_L(t) = A_L\cos[\omega_c t + \theta(t)] \quad (10)$$

and the time variation is represented by the phase term, $\theta(t)$. For cases when the signal amplitude is much larger than the noise components, the phase noise can be approximated as

$$\theta(t) \approx \frac{i_{ns}(t)}{(I_{ph}/2) + i_{nc}(t)} \approx \frac{i_{ns}(t)}{I_{ph}/2}. \quad (11)$$

The total mean square noise is given by

$$\overline{i_n^2} = \frac{\overline{i_{nc}^2}}{2}(1 + \cos 2\omega_c t) + \frac{\overline{i_{ns}^2}}{2}(1 - \cos 2\omega_c t) + 2\overline{i_{nc} i_{ns}}(\cos\omega_c t \sin\omega_c t). \quad (12)$$

Since the second harmonics are outside the passband of the circuit, and the quadrature noise components are independent, Equation 12 becomes

$$\overline{i_n^2} = \frac{1}{2}(\overline{i_{nc}^2} + \overline{i_{ns}^2}). \quad (13)$$

Assuming that $\overline{i_{ns}^2} = \overline{i_{nc}^2}$, the mean square phase noise from Equation 11 is

$$\overline{\theta_n^2} = \frac{\overline{i_{ns}^2}}{I_{ph}^2/4} = \frac{4\overline{i_n^2}}{I_{ph}^2}. \quad (14)$$

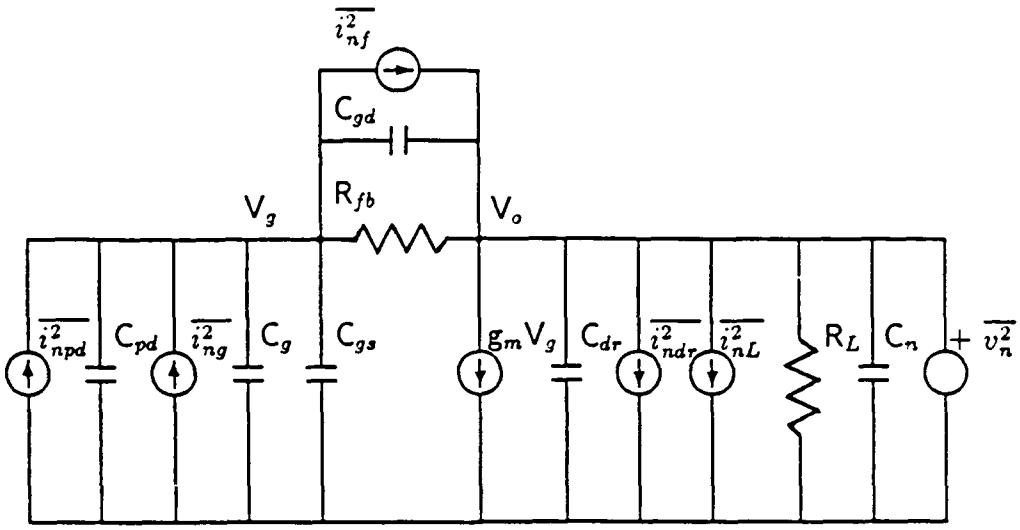

Figure 16: Noise circuit model for CMOS transimpedance amplifier

A circuit model for noise analysis of the first stage amplifier is shown in Figure 16. This model represents the small signal nature of the first stages of the receiver amplifier chain. In this model, the noise sources are independent of each other. $\overline{i_{npd}^2}$ is the mean square noise current at the photodetector. Noise due to the optical source such as relative intensity laser noise can be included here, and the combination of thermal diode junction noise and the source noise is given as

$$\overline{i_{npd}^2} = 4kTB/r_j + \overline{i_{n,source}^2} \tag{15}$$

where k is Boltzmann's constant, T is the temperature in Kelvin, B is the noise bandwidth of the receiver, and r_j is the reverse bias junction resistance of the photodiode. The shot noise portion of $\overline{i_{npd}^2}$ is neglected, as is the 1/f noise. These are good assumptions due to the reverse bias on the diode and the high frequency operation, respectively. $\overline{i_{ng}^2}$ is the shot noise due to leakage in the gate of the transistor, represented by

$$\overline{i_{ng}^2} = 2qI_gB \tag{16}$$

where q is the electron charge and I_g is the gate leakage current. $\overline{i_{nf}^2}$ is the thermal

noise due to the resistance of the feedback path, calculated as

$$\overline{i_{nf}^2} = 4kTB/R_{fb}. \tag{17}$$

$\overline{i_{ndr}^2}$ is the drain-source shot noise, represented by

$$\overline{i_{ndr}^2} = 4kT(2/3)g_mB \tag{18}$$

where g_m is the transconductance of the pull-down transistor for the first stage. Again, 1/f noise has been neglected in the calculation of the channel noise current. $\overline{i_{nL}^2}$ is the thermal noise due to the load resistance, given by

$$\overline{i_{nL}^2} = 4kTB/R_L. \tag{19}$$

The equivalent mean square noise voltage at the input to the second stage is determined by a similar analysis for noise sources in that stage, and is given as

$$\overline{v_n^2} = \frac{B\{2R_L^2qI_g + 4kT(2/3g_m)[1+\omega^2R_L^2(C_{dr}+C_g)^2]\}}{1+\omega^2(R_LC_{dr})^2} \tag{20}$$

This calculation assumes that the second stage dimensions are identical to those of the first stage; therefore, the load resistances, gate capacitances, and transconductances are equal.

For the noise analysis, the effects of each source on the output voltage of the first stage are determined, and an equivalent input current source for each is derived. Since $\overline{i_{npd}^2}$ and $\overline{i_{ng}^2}$ appear at the input, there is no need to calculate equivalent values for them. The equivalent input noise due to $\overline{i_{nf}^2}$ is

$$\overline{i_{neq,f}^2} = \overline{i_{nf}^2}\left|\frac{R_{fb}[(g_mR_L-\omega^2R_LC_a\tau_L)+j\omega(g_mR_L\tau_L-C_aR_L)]}{R_L[(1-g_mR_{fb}-\omega^2G\tau_f\tau_G)+j\omega(\tau_f+\tau_L-g_mR_{fb}\tau_L)]}\right|^2 \tag{21}$$

where $\tau_L = R_LC_{dr}$, $\tau_f = R_{fb}C_{gd}$, $G = (1+g_mR_L)$, and $C_a = (C_g + C_{gs} + C_{pd})$. The equivalent input noise due to $\overline{i_{ndr}^2}$ is

$$\overline{i_{neq,dr}^2} = \tag{22}$$
$$\overline{i_{ndr}^2}\left|\frac{R_L\{[1+2g_mR_L-\omega^2(G\tau_G\tau_f+R_{fb}C_a\tau_L)]+j\omega[G(\tau_G+\tau_f)+R_{fb}C_a]\}}{R_L[(1-g_mR_{fb}-\omega^2G\tau_f\tau_G)+j\omega(\tau_f+\tau_L-g_mR_{fb}\tau_L)]}\right|^2$$

where $\tau_G = \tau_L/G$. The equivalent input noise due to $\overline{i_{nL}^2}$ is similar, and given as

$$\overline{i_{neq,L}^2} = \overline{i_{nL}^2}\left|\frac{R_L\{[1+2g_mR_L-\omega^2(G\tau_G\tau_f+R_{fb}C_a\tau_L)]+j\omega[G(\tau_G+\tau_f)+R_{fb}C_a]\}}{R_L[(1-g_mR_{fb}-\omega^2G\tau_f\tau_G)+j\omega(\tau_f+\tau_L-g_mR_{fb}\tau_L)]}\right|^2. \tag{23}$$

The input equivalent noise due to the next stage is

$$\overline{i_{neq,n}^2} = \frac{\omega^2C_a^2(1+\omega^2\tau_f^2)}{1+(C_aR_{fb}+\tau_f)^2\omega^2}\overline{v_n^2}. \tag{24}$$

The total input mean square noise current is the sum of the independent equivalent input noise sources, or

$$\overline{i_{neq,t}^2} = \overline{i_{neq,pd}^2} + \overline{i_{neq,g}^2} + \overline{i_{neq,f}^2} + \overline{i_{neq,dr}^2} + \overline{i_{neq,L}^2} + \overline{i_{neq,pd}^2}. \tag{25}$$

The total equivalent input noise can then be compared to the mean square photocurrent, and an equivalent input SNR can be determined.

An important parameter in the quantification of the receiver noise is the input noise bandwidth of the transimpedance amplifier. In order to determine this bandwidth, the output voltage characteristic transfer function has been developed from circuit analysis. The transfer function is given as

$$v_o/i_{in} = N(s)/D(s) \tag{26}$$

where

$$N(s) = R_L[(1-g_mR_{fb})+(\tau_f+G\tau_G-g_mR_{fb}\tau_L)s+G\tau_f\tau_Gs^2] \tag{27}$$

and

$$\begin{aligned} D(s) \;=\; & G+[G(\tau_f+\tau_G+\tau_L)+(R_{fb}+R_L)C_a]s+ \\ & [G(\tau_f\tau_G+\tau_L(\tau_f+\tau_G))+2R_{fb}C_a\tau_L+R_LC_a(\tau_f+\tau_L)]s^2+ \\ & [G\tau_L\tau_f\tau_G+R_{fb}C_a\tau_L^2+R_LC_a\tau_f\tau_L]s^3. \end{aligned} \tag{28}$$

In order to determine typical noise performance for the transimpedance receiver, and to compare it with the performance of a phase-locked loop receiver

Symbol	Value	Symbol	Value
R_L	2kΩ	C_i	66.3 nF/cm^2
R_{fb}	20kΩ	C_g	2 pF
r_j	1MΩ	C_{gs}	0.52 pF
g_m	0.0134 Ω^{-1}	C_{gd}	0.54 pF
τ_f	10.8 ns	C_{dr}	4.3 pf
τ_L	12.6 ns	B	35 MHz

Table 1: Circuit element values.

found in Chapter 4, the input noise bandwidth and circuit model element values have been calculated, based on measured parameters from the MOSIS fabrication process for SPICE simulation input. These values are shown in Table 1, and represent a transimpedance amplifier implementation with 3 micron channel length transistors.

For the example circuit parameters given in Table 1, the passbands of the individual input equivalent noise sources are limited by the input noise passband. The input equivalent noise currents for the example can therefore be approximated as

$$\overline{i^2_{neq,f}} \approx \overline{i^2_{nf}} \tag{29}$$

$$\overline{i^2_{neq,dr}} \approx \overline{i^2_{ndr}}\left[\frac{2R_{fb}}{R_L}\right]^2 \tag{30}$$

$$\overline{i^2_{neq,L}} \approx \overline{i^2_{nL}}\left[\frac{2R_{fb}}{R_L}\right]^2 \tag{31}$$

$$\overline{i^2_{neq,n}} \approx \frac{\overline{v^2_n}}{R^2_{fb}}. \tag{32}$$

The respective contributions of the individual noise sources and the corresponding mean square equivalent input noise current for the circuit parameters shown in the example are listed in Table 2. The total equivalent input mean square noise current for this receiver can be approximated as 1.91×10^{-16} (A^2). The estimated mean square phase jitter can be determined from Equation 14. For the example given, this is 7.64×10^{-6}radians^2. The rms phase jitter is 2.76×10^{-3} radians.

Source	Mean Square Noise	Mean Square Eq. Input Noise
	(A^2 or V^2)	(A^2)
$\overline{i_{npd}^2}$	5.76×10^{-19}	5.76×10^{-19}
$\overline{i_{nq}^2}$	1.12×10^{-23}	1.12×10^{-23}
$\overline{i_{nf}^2}$	2.88×10^{-17}	2.88×10^{-17}
$\overline{i_{ndr}^2}$	5.15×10^{-15}	1.03×10^{-16}
$\overline{i_{nL}^2}$	2.88×10^{-16}	5.77×10^{-17}
$\overline{v_n^2}$	1.31×10^{-19}	3.26×10^{-28}

Table 2: Individual noise source values and equivalent input noise values for 3-micron example.

In order to compare the magnitude of the noise-induced timing uncertainty for signals of different frequencies, a conversion from rms phase jitter in units of radians to rms time jitter in units of seconds can be performed. If the transimpedance amplifier frequency were not limited by slew rate, but rather by the noise bandwidth of the receiver, maximum clock frequency for the example given would be around 35 MHz. This corresponds to an rms signal jitter of 12.5 ps. In reality, however, the maximum clock frequency is likely to be limited by the slew rate of the amplifier. In Chapter 5, laboratory measurements for test receivers implemented in 3-micron CMOS indicate that the maximum clock frequency for reasonable optical input power is around 15 MHz. The phase jitter for a 15 MHz clock signal based on the example circuit parameters corresponds to an rms signal jitter of 29.3 ps.

For the purpose of estimating the noise performance of an implementation of the transimpedance amplifier with 2 micron channel length transistors, the noise bandwidth is scaled by a factor of 1.5 to 52.5 MHz. The equivalent mean square noise current similarly scales to 2.86×10^{-18} A^2, and the mean square phase noise becomes 1.15×10^{-5} radians^2 with an rms value of 3.39×10^{-3} radians. For an operating frequency of 50 MHz, the skew-limited maximum operating frequency for a 2-micron design as simulated in Section 3.2, the rms signal jitter is estimated as 10 ps.

The noise contribution to timing uncertainty for the transimpedance amplifier that has been presented in this chapter is two orders of magnitude less than the static skew that is expected for the same circuit. For this reason, the jitter can be neglected in determining the maximum operating frequency for the clock. The skew contribution, however, is quite large. This receiver design provides no real synchronization improvement over metal wired distribution systems for chip sizes less than 1 cm, while requiring an inhibitively large amount of layout area for implementation of each receiver. While some application might exist for wafer scale clock distribution using a transimpedance amplifier approach, a receiver design with better performance and less layout overhead is required for single chip optical clock distribution. A phase-locked loop receiver with picosecond static response skew and sub-picosecond rms signal jitter is presented in Chapter 4 which is expected to require much less chip area and will operate for a 100-200 MHz clock signal.

Chapter 4

Phase-Locked Loop Receiver

The limitations of the performance of the transimpedance MOS receiver make it inappropriate for chip level optical clock distribution systems. The low gain, poor response time, large response skew, and massive layout area requirements indicate that an alternate receiver approach must be used if multiple receivers are to be implemented on a single chip with performance superior to existing metal clock distribution systems.

An attractive alternative receiver is the phase-locked loop receiver that is presented in this chapter. This receiver is based on a voltage controlled ring oscillator, which allows implementation with significantly lower transistor size requirements, while operating in the 100–200 MHz frequency range. A description of this circuit and its operation is presented in Section 4.1.

As with the transimpedance amplifier, the phase-locked loop receiver is subject to static and temporal signal transition uncertainties. Section 4.2 is a presentation of fabrication-related response skew, while Section 4.3 is a description of the signal jitter due to noise at each receiver. For each case, the phase-locked loop with the same optical input power represents orders of magnitude improvement over the transimpedance receiver, as is demonstrated by examples of a 2-micron

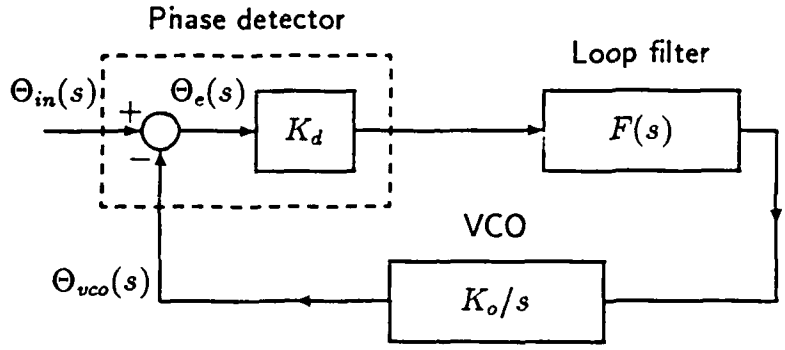

Figure 17: A phase-locked loop block diagram.

implementation of the phase-locked loop receiver and comparisons with the corresponding examples for the transimpedance receiver presented in Chapter 3. A comparison of incident optical power levels required for the two receiver designs to achieve an arbitrary minimum timing uncertainty is presented in Section 4.4.

4.1 Circuit Descriptions

Phase-locked loop receivers operate by matching the phase and frequency of a locally generated signal with an external waveform. Traditionally, a phase-locking circuit consists of a voltage controlled oscillator (VCO), a phase detector, and a loop filter, as shown in Figure 17.

The phase detection and filtering produce a control voltage which is used as an input to the voltage controlled oscillator and ultimately controls the output frequency and phase of the loop. The phase detector measures the separation of the input and VCO output signals. The phase error voltage is a function of the time separation between the corresponding pulse edges of the two signals. The two signals become phase-locked when the phase separation between the two signals becomes a contant value. By adjusting the control voltage for the VCO,

the output frequency and phase are adjusted until the separation between the signals reaches a steady-state value, bringing the VCO output signal into lock with the input signal.

There are typically three input frequency ranges that can be used to describe the phase aquisition and tracking ability of phase locked loops. These ranges are centered around the output frequency of the VCO. The first two ranges describe the dynamic capability of the loop to track or reacquire an input signal with modulation. The third describes the ability of the loop to lock to a single frequency.

The narrowest range, the lock-in range or $\Delta\omega_L$, describes the maximum input frequency shift for which the loop can maintain phase-lock between the input and VCO signals. Any frequency steps less than $\Delta\omega_L$ for the input signal can be followed by the VCO without losing phase lock. If the input frequency is modulated outside this range, the phase locking between the two signals can slip a few cycles, and perhaps be recovered. This is relevant to frequency or phase modulation receivers, because it defines the bandwidth of the input modulation frequency that can be tracked, but it has limited meaning for use with a single frequency clock signal input.

The second range, the pull-in range or $\Delta\omega_P$, describes input frequencies which can be pulled in to lock after slipping for a number of cycles. Any frequency steps in this range between the input frequency and the VCO output can cause the phase to slip several cycles, but the system will again return to a locked condition eventually. The time required to return to lock is dependent on the dynamics of the phase locked loop and the magnitude of the frequency step. Again, for clock operation, this parameter has limited meaning, because the input frequency for the clock signal will only be stepped when the clock signal is turned on.

The widest range, the hold range or $\Delta\omega_H$, describes the input frequency range in which a steady-state phase error exists. It is primarily defined by the saturation limits of the phase detector. Input frequencies that differ from the VCO free-running signal by more than $\pm\Delta\omega_H$ have an expected error voltage that excedes

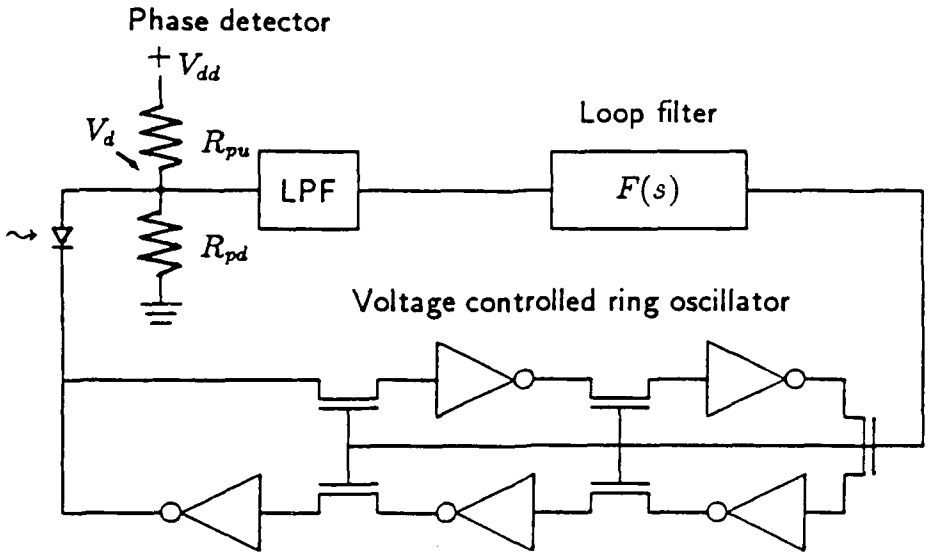

Figure 18: A phase-locked loop for optical clock detection.

the output voltage available from the phase detector, and therefore, the loop will not lock. The VCO can lock to a single frequency clock signal inside the hold range, however changes in the input frequency could result in instability in which the loop can become momentarily unlocked, if the changes are larger than $\Delta\omega_P$.

Phase-locking systems can be designed with either digital components, or as analog circuits. Digital phase-locked loops have the advantage of programmability and digital level output signals, but operate at relatively low frequencies. Analog phase-locked loops can operate at higher frequencies, but typically require large devices for implementing the VCO and loop filters, while usually delivering a sinusoidal small signal output.

The receiver presented here is a hybrid of analog and digital approaches, modeled after those recently reported in the literature[21,27]. Shown in Figure 18, this circuit has analog RC low pass filters, an analog phase detector, and a voltage controlled ring oscillator. The ring oscillator provides a digital output level signal which can then be used as an input for the two-phase clock synthesizing circuit described in Section 3.1.2, while being implemented with minimum size

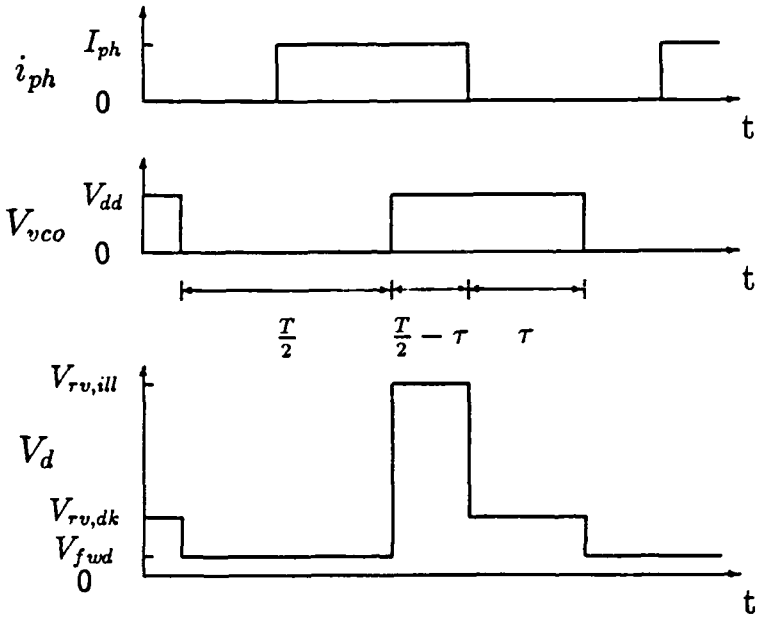

Figure 19: Timing diagram of input signal, VCO output signal, and photodiode output for phase detection.

transistors.

The design reported by Fried [21] has been modified by actively biasing a photodiode with the VCO output signal to allow photodetection and phase comparison together, eliminating the need for a separate signal multiplying device. When the output voltage of the VCO is high, the photodiode is reverse biased, and a photocurrent flows when light is present on the diode. When the output voltage of the VCO is low, the photodiode is forward biased, and any photocurrent is negligible when compared to the forward diode current. In this manner, V_d, the diode voltage indicated in Figure 18, is modulated between three states, as shown in Figure 19. The voltage levels for the three states of V_d are: V_{fwd}, the forward biased junction voltage for the diode, $V_{rv,dk}$, the reverse biased junction voltage with no illumination, and $V_{rv,ill}$, the reverse biased junction voltage with

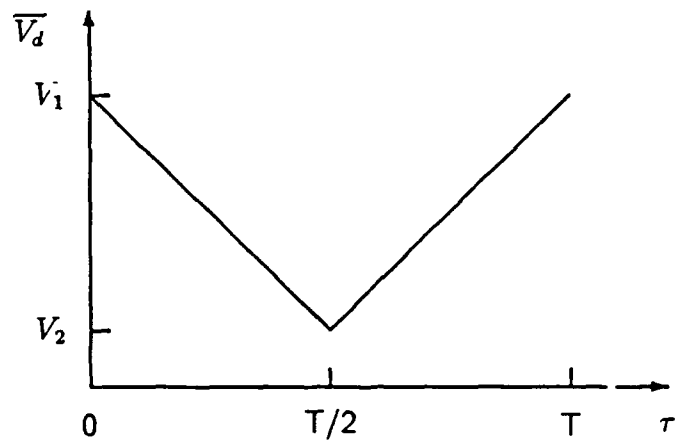

Figure 20: Phase detector conversion characteristic as a function of signal separation, τ.

illumination, where

$$V_{fwd} = \frac{kT}{q} \ln(I_f/I_o) \tag{33}$$

$$V_{rv,dk} = \frac{V_{dd}R_{pd}}{R_{pu} + R_{pd}} \tag{34}$$

and

$$V_{rv,ill} = \frac{V_{dd}R_{dd} + I_{ph}R_{pu}R_{pd}}{R_{pu} + R_{pd}}. \tag{35}$$

The phase error signal for the loop in this phase-locked loop is the DC term of the output signal from the photodiode, and is produced by low pass filtering V_d.

The characteristic phase detector output voltage from the low pass filter is shown in Figure 20, and is the time average of the voltage on the photodiode, given as

$$\overline{V_{out}} = \begin{cases} \frac{1}{2}\left[\frac{V_{dd}R_{pd}}{R_{pu}+R_{pd}} + \frac{I_{ph}R_{pu}R_{pd}}{R_{pu}+Rpd} + V_{fwd}\right] - \frac{\tau I_{ph}R_{pu}R_{pd}}{T(R_{pu}+R_{pd})} & \text{for } 0 \le \tau \le \frac{T}{2} \\ \frac{1}{2}\left[\frac{V_{dd}R_{pd}}{R_{pu}+R_{pd}} - \frac{I_{ph}R_{pu}R_{pd}}{R_{pu}+Rpd} + V_{fwd}\right] + \frac{\tau I_{ph}R_{pu}R_{pd}}{T(R_{pu}+R_{pd})} & \text{for } \frac{T}{2} \le \tau \le T \end{cases} \tag{36}$$

where τ is the time separation between the VCO output pulse edge and the optical input signal pulse edge. For a square wave optical input signal, the characteristic

function of the phase detector is linear, as shown, and the conversion gain, K_d, is given as the slope of this characteristic, or

$$|K_d| = \frac{I_{ph}(R_{pu}||R_{pd})}{(T/2)} = I_{ph}R_{eq}/\pi \tag{37}$$

where R_{pu} and R_{pd} are the pull-up and pull-down resistors in a voltage divider, R_{eq} is the equivalent parallel resistance of R_{pu} and R_{pd}, and T is the period of the VCO output signal. The maximum and minimum values of $\overline{V_d}$ are given as

$$V_1 = \frac{1}{2}\left[\frac{V_{dd}R_{pd} + V_{dd}R_{pu}R_{pd}}{R_{pu} + R_{pd}} + V_{fwd}\right] \tag{38}$$

and

$$V_2 = \frac{1}{2}\left[\frac{V_{dd}R_{pd}}{R_{pu} + R_{pd}} + V_{fwd}\right] \tag{39}$$

respectively.

The operation of low order phase-locked loops is well understood [35,36]. In general, the order of a phase-locked loop is determined by the Laplace transform of the phase transfer function for the loop. Since the phase of the VCO output signal can be represented as the integral of the frequency of that signal, the transfer function for the VCO block in the loop can be represented as

$$H_{vco}(s) = K_o/s \tag{40}$$

where K_o is the voltage-to-frequency conversion gain of the VCO. If the phase-to-voltage conversion gain of the phase detector is K_d, as defined in Equation 37, and the transfer function of the loop filter is $F(s)$, then the loop phase transfer function can be represented as

$$H(s) = \frac{\Theta_{vco}(s)}{\Theta_{in}(s)} = \frac{K_oK_dF(s)}{s + K_oK_dF(s)}. \tag{41}$$

While simpler phase-locked loops could be used for clock receiver designs, perhaps the best is a second order loop with a lag-lead low pass filter, shown in Figure 21. This design allows a wide frequency range for operation, has independent noise bandwidth and loop gain, is easily implemented with reasonable layout requirements, and has a small steady-state phase error.

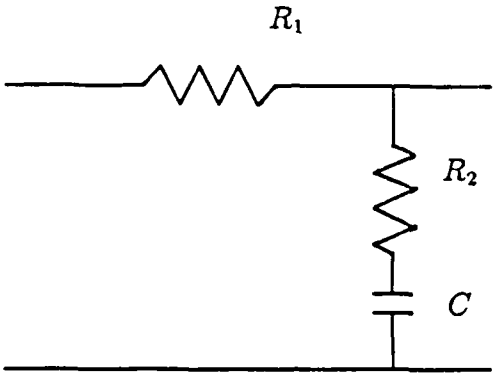

Figure 21: Lag-lead low pass loop filter.

The tranfer function of the lag-lead low pass loop filter is given as

$$F(s) = \frac{s\tau_2 + 1}{s\tau_1 + 1} \tag{42}$$

where $\tau_1 = (R_1 + R_2)C$ and $\tau_2 = R_1C$. The transfer function for the phase-locked loop with this type of filter can be represented as [37]

$$H(s) = \frac{K_oK_d(s\tau_2 + 1)/\tau_1}{s^2 + s(1 + K_oK_d\tau_2/\tau_1) + K_oK_d/\tau_1}. \tag{43}$$

Typically the transfer function is represented in a canonical form for second order circuits, or

$$H(s) = \frac{s[2\zeta\omega_n - \omega_n^2/(K_oK_d)] + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \tag{44}$$

where ω_n is the natural frequency of the second order system, and ζ is a damping factor. By matching the coefficients of Equation 43 and Equation 44, ω_n and ζ can be expressed in terms of K_d, K_o, τ_1 and τ_2, yielding

$$\omega_n = \sqrt{K_oK_d/\tau_1} \tag{45}$$

$$\zeta = \frac{1}{2}\sqrt{\frac{K_oK_d}{\tau_1}}\left(\tau_2 + \frac{1}{K_oK_d}\right). \tag{46}$$

Most of the interesting operating parameters of the phase-locked loop can be expressed in terms of ω_n, ζ, and the conversion gains.

For the noise analysis that is presented in Section 4.3, the loop bandwidth, B_L, is of particular interest. The noise bandwidth is defined as [38]

$$B_L = \int_0^\infty |H(j2\pi f)|^2 df \tag{47}$$

where $H(j2\pi f)$ is the one-sided transfer function of the loop to be consistent with Gardner's phase-locked loop terminology [38]. For a second order loop with a lag-lead passive low pass filter, is

$$B_L = \frac{\left(\frac{K_o K_d \tau_2}{\tau_1} + \frac{1}{\tau_2}\right)}{4\left(1 + \frac{1}{K_o K_d \tau_2}\right)} \qquad \text{Hz.} \tag{48}$$

The noise bandwidth has dimensions of hertz rather than radians/second to be compatible with Gardner's terminology in which noise power spectral density is expressed in units of watts/hertz.

In addition to the noise bandwidth of the loop, the phase acquisition frequency ranges that have been described be determined[39]. For the second order loop with a passive lag-lead filter, the hold range is

$$\Delta\omega_H = K_o K_d F(0) = K_o K_d. \tag{49}$$

This is the frequency parameter of primary concern for clock receiver designs; however, for means of comparison, the pull-in range for the same system with sinusoidal phase detector characteristics when $K_o K_d \gg \omega_n$, can be expressed as

$$\Delta\omega_P \approx \frac{8}{\pi}\sqrt{\zeta\omega_n K_o K_d - \omega_n^2/2}. \tag{50}$$

The lock-in range for second order loop with a sinusoidal phase detector can be approximated as the natural frequency of the loop, or

$$\Delta\omega_L \approx \omega_n. \tag{51}$$

Since the phase detector characteristic for the optical clock receiver is triangular, rather than sinusoidal, the lock-in and pull-in ranges are extended by a factor of $\pi/2$ [40].

Parameter	Value	Parameter	Value
R_1	32.3 kΩ	C	0.5 pF
R_2	31.4 kΩ	τ_1	31.4 ns
R_{eq}	31.4 kΩ	τ_2	15.3 ns

Table 3: Typical phase-locked loop filter and phase detector parameters.

Two time parameters predicting the acquision process are T_P, the time to pull the phase to within 2π of the input signal, and T_L, the time for the phase to lock, once the frequency is pulled into the lock range. As might be expected, the pull-in time is generally much longer than the lock-in time. For second order loops [39]

$$T_P \approx \frac{\Delta\omega_o^2}{2\zeta\omega_n^3} \tag{52}$$

where $\Delta\omega_o$ is the difference between the input frequency and the center frequency of the VCO. The lock-in time for second order systems can be approximated as

$$T_L \approx 1/\omega_n. \tag{53}$$

In order to estimate the operating parameters of a typical implementation of the voltage controlled ring oscillator with minimum size inverter gates and 2 micron channel lengths, SPICE simulations have been performed to determine the voltage-to-frequency conversion gain. In this study, the control voltage for the pass transistors was varied and the simulated VCO frequency was determined. The resulting VCO frequency characteristic is shown in Figure 22. For this investigation, NMOS pass transistors of minimum size were used, allowing the highest operating frequencies for the VCO. The result was linearized to obtain a conversion gain of $K_o = 2\pi(128)\times10^6$ radians/V-s.

Parameters for the phase detector and loop low pass filter that might be considered typical for a 2-micron implementation of this phase-locked loop are shown in Table 3. These estimates reflect the values that might be implemented with reasonable layout area requirements.

Given the filter parameters shown in Table 3 and assuming a photocurrent of 10 μA, an estimate of typical operating parameters for the phase-locked loop can be

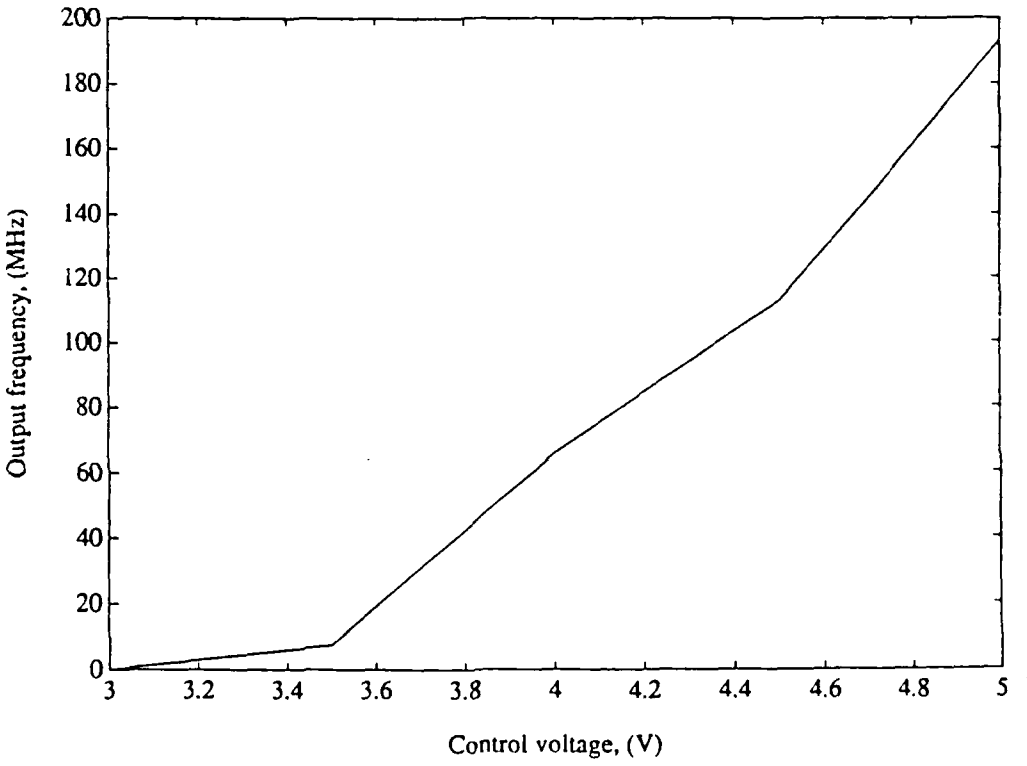

Figure 22: VCO characteristic curve for implementation with 2-micron channel length transistors.

Parameter	Symbol	Value
VCO conversion gain	K_o	804×10^6 radians/s-V
Phase detector conversion gain	K_d	0.1 V/radian
Loop natural frequency	ω_n	50.6×10^6 radians/s
Loop damping factor	ζ	0.7
Noise bandwidth	B_L	14.4 MHz
Hold range	$\Delta\omega_H$	80.4×10^6 radians/s (12.8 MHz)
Pull-in range	$\Delta\omega_P$	80.4×10^6 radians/s (12.8 MHz)
Lock-in range	$\Delta\omega_L$	79.5×10^6 radians/s (12.7 MHz)
Pull-in time	T_P	35.6 ns
Lock-in time	T_L	19.8 ns

Table 4: Typical phase-locked loop parameter for a 2-micron CMOS implementation.

determined. These are shown in Table 4. The phase detector conversion gain, K_d, given the expected photocurrent and equivalent input resistance, is 0.1 V/radian. As indicated in the table, the natural frequency of the second order loop, ω_n, is 50.6×10^6 radians/s, and the damping factor, ζ, is 0.7. In this design, ζ has been optimized by adjusting the pole and zero of the loop filter to allow minimum decay time for the transient response of the loop. The noise bandwidth for the phase-locked loop, from Equation 48, is 14.4 MHz. Using the estimated values of ω_n and ζ, the estimated hold range for the phase-locked loop is 80.4×10^6 radians/s, or 12.7 MHz. Since the relationship $K_oK_d \gg \omega_n$ is not true for this example, Equation 50 is not valid, and the pull-in range is equal to the hold range. While the lock-in range is not a critical parameter for clock syncronization, assuming a triangular phase detector characteristic, it can be estimated as 79.5×10^6 radians/s, or 12.7 MHz. The pull-in time and lock-in time can be estimated as 35.6 ns and 19.8 ns, respectively, yielding a total expected acquisition time of 55.4 ns for an input frequency of maximum separation from the free-running oscillator frequency.

The phase-locked loop receiver that has been given in this example demonstrates two improvements over the transimpedance receiver design presented in Chapter 3. The operating frequency of the phase-locked loop is much higher than the maximum available with the transimpedance receiver with similar transistor

channel lengths. The phase-locked loop can operate in the frequency range of 100 – 200 MHz, while the transimpedance receiver is limited to about 50 MHz. In addition, the digital nature of the voltage controlled ring oscillator in the phase-locked loop allows implementation with minimum size transistors. The largest devices in the phase-locked loop are the capacitors used for the low pass filters; these can be realized using the gate capacitance of transistors, requiring 1/4 of the area required for a single input inverter in the transimpedance design. Overall, the phase-locked loop can be designed in a layout area much smaller than the transimpedance receiver.

4.2 Static Response Skew

Just as has been shown for the transimpedance clock receiver presented in Chapter 3, identically drawn circuit elements for the phase-locked loop receiver cannot be fabricated without small variations in dimensional and threshold parameters. These parameter variations between copies of phase-locked loops implemented on a given VLSI chip can cause a variation in the steady-state phase error from one receiver to another. This difference in steady-state phase error represents the clock skew in the phase-locked loop receiver approach for optical clock distribution, and is discussed in this section.

When a phase-locked loop reaches a locked state, a steady-state difference between the input and oscillator signal phases is maintained. This steady-state phase error is determined by using the final value theorem of Laplace transforms [41]. The Laplace transform of the phase error, $\Theta_e(s)$, can be represented as

$$\Theta_e(s) = \frac{s\Theta_i(s)}{s + K_o K_d F(s)}. \tag{54}$$

From the final value theorem, the steady-state time response can be determined by the relationship

$$\lim_{t\to\infty} y(t) = \lim_{s\to 0} sY(s). \tag{55}$$

Applying the final value theorem to the phase error, the steady-state phase error, θ_{ss}, is given as

$$\theta_{ss} = \lim_{s \to 0} \frac{s^2 \Theta_i(s)}{s + K_o K_d F(s)}. \tag{56}$$

To model the steady state error when a single frequency input is applied to an unlocked loop, the input phase is represented as a step in frequency, $\Delta\omega_o/s^2$, where $\Delta\omega_o$ is the frequency difference between the input signal and the free-running local oscillator. Substituting this value of $\Theta_i(s)$ in Equation 56. the steady state phase error is represented as

$$\theta_{ss} = \lim_{s \to 0} \frac{\Delta\omega_c}{s + K_o K_d F(s)} = \frac{\Delta\omega_o}{K_o K_d F(0)}. \tag{57}$$

For phase-locked loops with passive loop filters, $F(0) = 1$, and Equation 57 becomes

$$\theta_{ss} = \frac{\Delta\omega_o}{K_o K_d}. \tag{58}$$

Variations in the device dimensions and thresholds due to the fabrication process cause the conversion gains of the VCO and the phase detector to vary from one receiver to another on a chip, and consequently cause differences in steady-state phase error. The maximum difference in steady-state phase error due to device parameter variations can be represented as

$$\Delta\theta_{ss} = \theta_{max} - \theta_{min} = \Delta\omega_o \left(\frac{1}{(K_o K_d)_{min}} - \frac{1}{(K_o K_d)_{max}} \right) \tag{59}$$

where θ_{max} and θ_{min} represent the steady-state phase errors for copies of the phase-locked loop with the maximum and minimum cases. For a worst-case estimate, $\Delta\omega_o$ represents the maximum frequency difference between the local input signal and the free-running oscillator, or $\Delta\omega_H$.

For comparison of the static synchronization uncertainty for the phase-locked loop receiver that has been presented in Section 4.1, SPICE simulations have been performed to determine the range of conversion gains expected for a phase-locked loop with 2 micron channel length transistors. The transistor channel dimensions were varied $\pm 0.1\mu$m, and the threshold voltages were varied ± 1 mV, as was done

for the transimpedance receiver analysis in Section 3.2. The minimum combination of conversion gains, $(K_o K_d)_{min}$, was computed as 78×10^6; the maximum value, $(K_o K_d)_{max}$, was 82×10^6. Using these values in Equation 59, along with the estimate of the hold range, $\Delta\omega_H$ =80.4×10^6, a typical value of the steady-state phase error is 0.05 radians. For an operating frequency of 150 MHz, this corresponds to a timing skew of 53 ps. When compared to the transimpedance receiver skew of 3 ns, the phase-locked loop static timing uncertainty represents an improvement of almost 2 orders of magnitude.

4.3 Signal Jitter Analysis

The phase-locked loop receiver approach has been shown in Section 4.1 to operate at a higher frequency than is possible with a transimpedance receiver, with static receiver skew that is orders of magnitude less than that obtained with the transimpedance approach, as shown in Section 4.2. In this section, a circuit noise analysis is presented, followed by a timing jitter analysis for the phase-locked loop with a comparison of clock jitter for the two receiver approaches.

The primary sources of noise in the phase-locked loop receiver are the photodiode and the biasing resistors. The thermal noise of the pass transistors and the inverter stages of the ring oscillator output signal is negligible due to the conversion gains of the VCO and the phase detector.

Since the photodiode is alternately forward and reverse biased by the output signal of the ring oscillator, the noise characteristic of the diode is not stationary. By separating the analysis into four cases of biasing and optical input signal levels, however, the noise characteristic for the photodiode can be approximated as stationary, and linear analysis can be performed. The four cases are summarized in Table 5. In general, a given cycle of the VCO output signal will cause the photodiode characteristic to have time segments representing each case, depending on the overlap between the optical input signal and the electronic VCO output signal. Since the noise contributions from each case are different, the total noise

Case	Description
Case I	VCO output LOW, optical input OFF
Case II	VCO output LOW, optical input ON
Case III	VCO output HIGH, optical input OFF
Case IV	VCO output HIGH, optical input ON

Table 5: Photodiode/phase detector states for noise analysis.

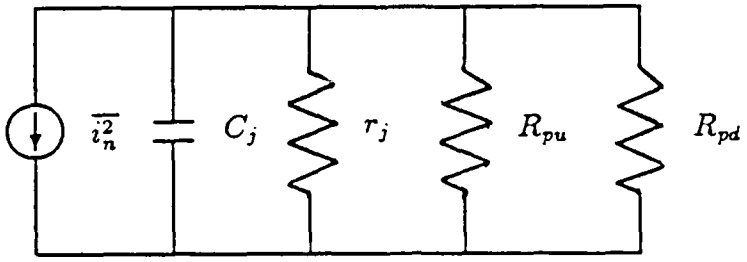

Figure 23: Noise model for phase-locked loop photodiode.

energy in a given cycle is a function of the overlap of the two signals. In the analysis presented here, the total noise energy for a given cycle is given first in terms of the overlap between signals, then is shown to be a weakly varying function of this overlap, so that the noise can be approximated as stationary. Since 1/f noise is not considered for this analysis due to the high operating frequency of the diode, the noise will also be assumed to have uniform spectral density.

The noise in the photodiode circuit can be modeled as a current source in parallel with the diode junction resistance and capacitance, as shown in Figure 23. In this model, $\overline{i_n^2}$ is given by [42]

$$\overline{i_n^2} = 4kTB/R_{eq} - 2qBI_d \tag{60}$$

where R_{eq} is the parallel combination of r_j, R_{pu} and R_{pd}, and B is the bandwidth of the circuit.

For cases *I* and *II*, the photodiode is forward biased. Since the photocurrent in *II* is much smaller than the forward current in the diode, I_d for both cases is approximately the forward current in the diode, I_f. Because the junction resistance for the forward biased diode is small, the equivalent resistance, R_{eq} can be

approximated as r_j, where

$$r_j = \frac{kT}{qI_f}. \tag{61}$$

Substituting Equation 61 into Equation 60, the total noise current for the forward biased cases is

$$\overline{i_{n,f}^2} = 4kTB/R_{eq} - 2kTB/R_{eq} = 2kTB/R_{eq}. \tag{62}$$

Assuming a square wave VCO output signal of approximately 50 per cent duty-cycle, the duration of the forward biased cases for a given cycle is

$$t_{fwd} = 1/(2f_{vco}) \tag{63}$$

where f_{vco} is the frequency of the VCO output signal. Combining Equation 62 and Equation 63, the total noise energy in a given cycle due to the forward biased cases is represented as

$$E_{n,fwd} = \frac{2kTB}{2f_{vco}}. \tag{64}$$

The noise current for the reverse biased cases (*III* and *IV*), can also be represented in terms of Equation 60. For these cases, the diode current is given as

$$I_{d,rev} = i_{ph} + I_{bg} + I_{dk} \tag{65}$$

where I_{bg} and I_{dk} are the photocurrent due to background light and the dark leakage current of the diode, respectively, and

$$i_{ph} = \begin{cases} 0 & \text{for Case } III \\ I_{ph} & \text{for Case } IV \end{cases}. \tag{66}$$

Combining Equation 60 and Equation 65, the total mean square noise current for the reverse biased cases is

$$\overline{i_{n,r}^2} = 4kTB/R_{eq} - 2qB(i_{ph} + I_{bg} + I_{dk}). \tag{67}$$

Since the junction resistance for the reverse biased diode is very large, the equivalent resistance for Equation 67 is the parallel combination of R_{pu} and R_{pd}.

For photocurrent levels of 10 μA or less, the shot noise contribution of Equation 67 is several orders of magnitude less than the thermal noise contribution, and the total equivalent mean square noise current for cases *III* and *IV* can be represented as

$$\overline{i_{n,r}^2} = 4kTB/R_{eq}. \tag{68}$$

The duration of the reverse biased cases for a given cycle of the VCO output signal is

$$t_{rev} = 1/(2f_{vco}) \tag{69}$$

The total noise energy for a given cycle of the VCO output signal due to the reverse biased cases can therefore be given as

$$E_{n,rev} = \frac{4kTB}{2f_{vco}}. \tag{70}$$

The total equivalent input noise power for phase-locked loop is the sum of the noise energy contributions from the forward and reverse biased cases divided by the period of the cycle, or

$$P_{in} = \frac{kTB/f_{vco} + 2kTB/f_{vco}}{1/f_{vco}} = 3kTB. \tag{71}$$

The noise power spectral density is

$$W_i = 3kT \tag{72}$$

Given the equivalent input noise power spectral density from Equation 72 and the noise bandwidth of the phase-locked loop from Section 4.1, the mean square phase noise on the output of the VCO can be represented as [38]

$$\overline{\theta_{no}^2} = \frac{W_i B_L}{P_s} \tag{73}$$

where P_s is the photocurrent input signal power, or

$$P_s = \frac{I_{ph}^2(R_{pu}||R_{pd})}{2} \tag{74}$$

where $R_{pu}||R_{pd}$ represents the equivalent input resistance and is the parallel combination of the pull-up and pull-down resistors in the voltage divider.

For the example set of device parameters presented in Section 4.1 including an equivalent input resistance of 31.4 kΩ, a noise bandwidth of 14.4 MHz, and an input photocurrent of 10 μA, the mean square phase noise is equal to 1.14×10^{-7} radians^2 and the rms phase noise is 3.37×10^{-4} radians. For an operating frequency of 150 MHz, this corresponds to an rms timing jitter of 0.36 ps. As is the case for the transimpedance receiver presented in Chapter 3, the timing uncertainty due to noise at each receiver is negligible when compared to the static response skew, but the overall noise performance of the phase-locked loop is orders of magnitude better than for the transimpedance receiver.

4.4 Optical Power Comparison

The difference between the timing uncertainy performance of the transimpedance receiver and the phase-locked loop receiver can be viewed in terms of the incident optical power required to achieve a given level of performance. Since the phase-locked loop receiver requires only that the optical signal synchronize existing electronic local oscillilators, the incident optical power for this receiver is much less than that required by the transimpedance receiver to achieve the same quality of timing performance. An analysis of the relationship between incident optical power and steady state response skew for both receivers is presented in this section, followed by a similar analysis for signal jitter.

The response skew for the transimpedance receiver is not related to the input signal strength, but is entirely dependent on the differences in transistor channel dimensions and threshold voltages between copies of the receiver on a given chip. For this reason, the response skew of the transimpedance receiver is constant with respect to input optical power.

The relationship between incident optical power and steady-state response

skew for the phase-locked loop receiver is a consequence of the input signal level dependency of the phase detector gain. Using K_d as defined by Equation 37 in Section 4.1, with the receiver response skew described by Equation 59 in Section 4.2, the relationship between the peak photocurrent and phase-locked loop response skew for an operating frequency, f, is

$$t_{sk,pll} = \frac{\Delta\omega_o \pi}{2\pi f I_{ph}} \left[\frac{1}{(K_o R_{eq})_{min}} - \frac{1}{(K_o R_{eq})_{max}} \right]. \tag{75}$$

The photocurrent is given by [43]

$$I_{ph} = \frac{P_{opt} \eta_{ext} q}{h\nu} \tag{76}$$

where P_{opt} is the incident optical power, η_{ext} is the external quantum efficiency, q is the charge on a electron, and $h\nu$ is the photon energy. Substituting Equation 76 for I_{ph} in Equation 75, the relationship between the response skew and incident optical power is given by

$$t_{sk,pll} = \frac{\Delta\omega_o \pi h\nu}{2\pi f \eta_{ext} q P_{opt}} \left[\frac{1}{(K_o R_{eq})_{min}} - \frac{1}{(K_o R_{eq})_{max}} \right]. \tag{77}$$

The external quantum efficiency includes reflection loss due to the difference in dielectric constants between silicon and air and is given as

$$\eta_{ext} = \eta_{int}(1 - R) \tag{78}$$

where η_{int} is the internal quantum efficiency and R is the power reflection coefficient at the silicon surface. For photodetection in silicon with illumination at a wavelength of 633 nm and a depletion layer width of 2 microns, the internal quantum efficiency is about 0.5 [44]. Assuming no antireflection coating, the reflection coefficient for the silicon surface is about 0.3. The external quantum efficiency, η_{ext}, is therefore approximately 0.35 for the photodiodes used in the receiver design examples presented in Chapter 3 and this chapter. $h\nu/q$ for 633 nm illumination is approximately 2 V^{-1}. Based on the example presented in Section 4.2, $(K_o R_{eq})_{min} = 2.52 \times 10^{13}$ s^{-1} and $(K_o R_{eq})_{max} = 2.57 \times 10^{13}$ s^{-1}. Using these values in Equation 77, assuming the maximum input frequency step,

$\Delta\omega_o = 80.4 \times 10^6$ radians/s and an operating frequency of 150 MHz, the response skew for the example phase-locked loop receiver is

$$t_{sk,pll} = \frac{1.18 \times 10^{-15}}{P_{opt}} \quad \text{(s)} \tag{79}$$

where P_{opt} is expressed in watts.

Figure 24 is a comparison of the response skew for the transimpedance and phase-locked loop receivers as a function of incident optical power. The response skew for the transimpedance amplifier is constant with respect to optical signal power, and is shown as the horizontal dashed line. The solid line represents the phase-locked loop response skew as given by Equation 79. As shown in the figure, the phase-locked loop can operate with much lower incident optical power to achieve a maximum skew less than the transimpedance receiver value.

In addition to comparing the response skew for the two receiver examples, the signal jitter due to noise can be compared. For each receiver, the rms signal jitter is a function of the incident optical power. The relationship of signal jitter to incident optical power for the transimpedance receiver is presented first, followed by a similar analysis for the phase-locked loop receiver.

From Equation 11, the transimpedance receiver rms signal jitter can be represented as

$$t_{ti,rms} = \frac{2\sqrt{\overline{i_n^2}}}{2\pi f I_{ph}} \tag{80}$$

where $\overline{i_n^2}$ is the mean square noise current, f is the operating frequency, and I_{ph} is the photocurrent. Substituting Equation 76 for I_{ph} in Equation 80, the rms signal jitter as a function of optical input power becomes

$$t_{ti,rms} = \frac{2\sqrt{\overline{i_n^2}}h\nu}{2\pi f P_{opt}\eta_{ext}q}. \tag{81}$$

Assuming an operating frequency of 50 MHz (the maximum frequency for a 2-micron implementation), $\overline{i_n^2} = 2.86\times10^{-18}$ A^2, $\eta_{ext} = 0.35$ and $h\nu/q = 2$ V^{-1}, the rms signal jitter for a 2-micron implementation of the transimpedance receiver is

$$t_{ti,rms} = \frac{7.17 \times 10^{-18}}{P_{opt}} \quad \text{(s)}\,. \tag{82}$$

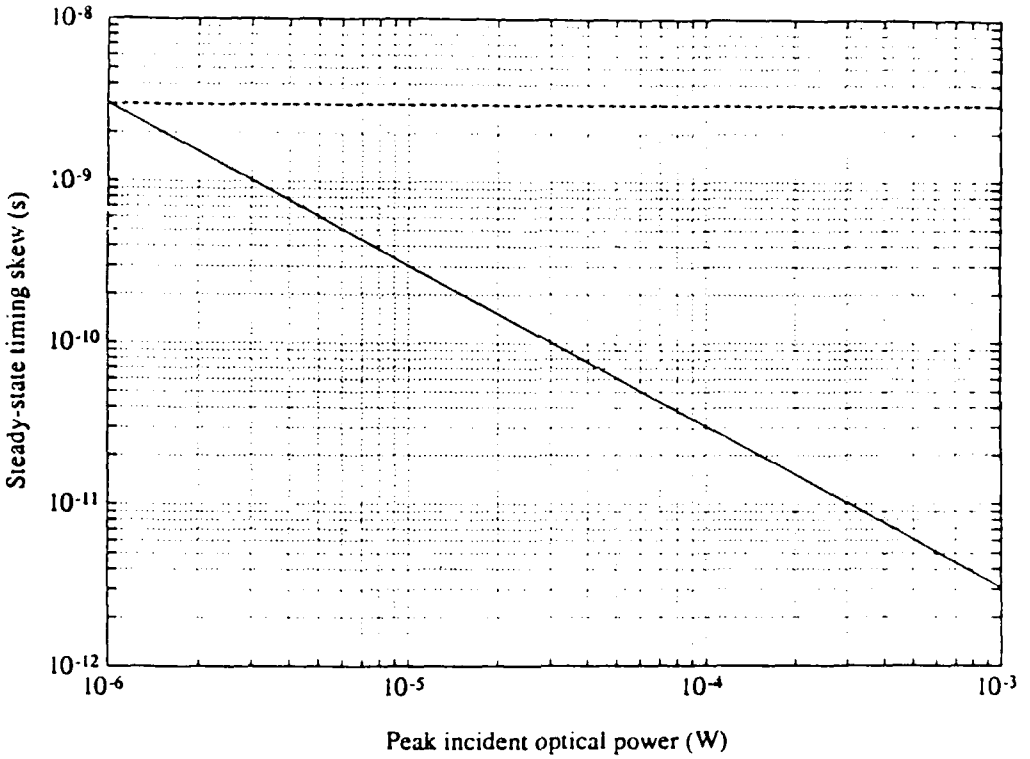

Figure 24: Response skew for transimpedance (dashed) and phase-locked loop (solid) receiver examples vs. peak incident optical power.

Using Equation 73 for the mean square phase noise for the phase-locked loop receiver, Equation 48 for the noise bandwidth, and Equation 76 for the photocurrent, the rms signal jitter for the phase-locked loop receiver can be represented as

$$t_{pll,rms} = \frac{1}{2\pi f}\left[\frac{W_i\left(\frac{\eta_{ext}P_{optq}R_{eq}K_o\tau_2}{h\nu\pi\tau_1}+\frac{1}{\tau_2}\right)}{2\left(\frac{\eta_{ext}P_{optq}}{h\nu}\right)^2 R_{eq}\left(\frac{\pi h\nu}{K_o R_{eq}\eta_{ext}P_{optq}\tau_2}+1\right)}\right]^{1/2} \tag{83}$$

where W_i is the equivalent noise power spectral density, and τ_1 and τ_2 are the time constants of the lead-lag low pass filter as defined in Section 4.1.

Using the example parameters given in Section 4.3 for the noise power spectral density and the phase-locked loop example parameters from Section 4.1, the rms signal jitter for a 2-micron implementation of the phase-locked loop receiver is

$$t_{pll,rms} \approx 2.69 \times 10^{-21}(6.85 \times 10^{11}P_{opt} + 6.53 \times 10^{7})^{1/2} \quad \text{(s)}. \tag{84}$$

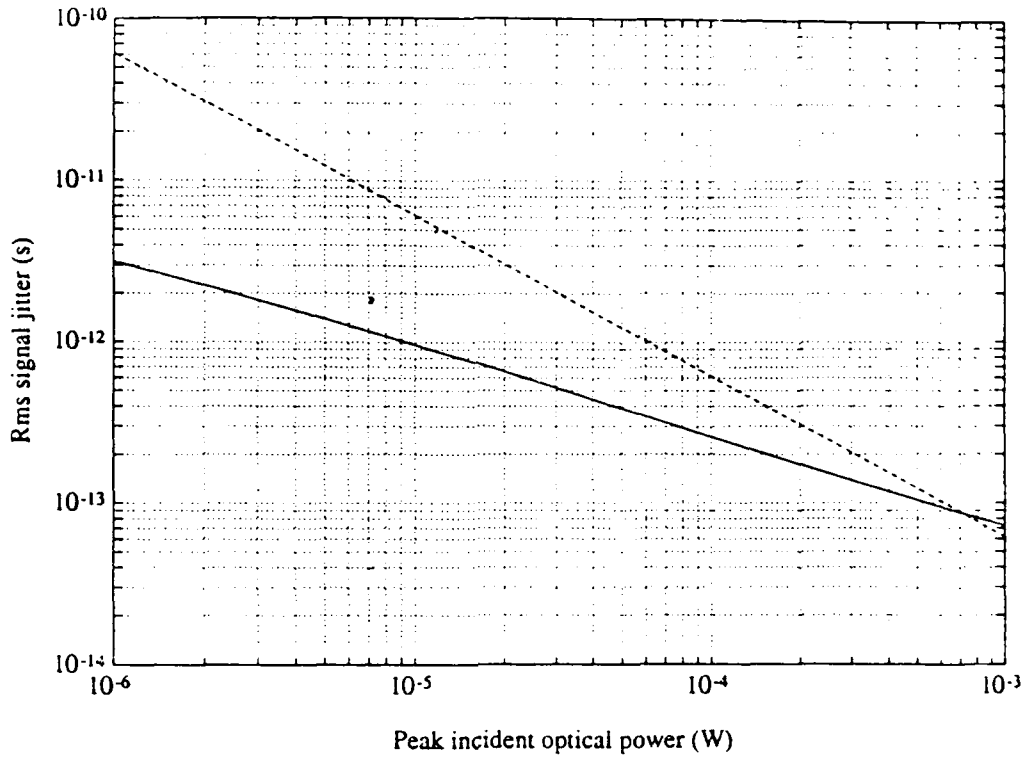

Figure 25: Rms signal jitter for transimpedance (dashed) and phase-locked loop (solid) receiver examples vs. peak incident optical power.

The rms signal jitter plots for the two receiver examples are shown in Figure 25 as a function of incident optical power. In the figure, the dashed line represents the transimpedance amplifier example, while the solid line represents the phase-locked loop receiver. Again, the phase-locked loop receiver can achieve a given rms signal jitter specification with a much lower incident optical power requirement.

A comparison of optical power costs, and operating parameters for the two receiver examples that have been presented clearly indicates that a phase-locked loop receiver is much more attractive for use in an optical clock distribution system for a VLSI chip. The phase-locked loop can be implemented with considerably less layout area and will operate at frequencies that represent a substantial improvement over wired clock distribution systems, where the transimpedance amplifier can offer at best marginal frequency improvement over wired distribution systems

while requiring massive use of chip surface area for implementation. In addition, the fabrication-related response skew between phase-locked loop receivers on a given chip is orders of magnitude less than the similar skew between transimpedance receivers for identical incident optical power levels. While the clock signal jitter for each type of receiver is small compared to the response skew, the phase-locked loop receiver also offers orders of magnitude improvement for this parameter as well. Timing uncertainty specifications can be maintainted with much lower optical power requirements for the phase-locked loop receiver example.

Chapter 5

Test Chip and Laboratory Measurements

A test chip has been designed to measure various operating parameters of a simple prototype transimpedance receiver and driver circuit for an optical clock distribution system. This chip contains circuits designed for measurement of the operating frequency and response skew for simple transimpedance optical receivers as well a set of circuits for measurement of leakage currents due to photogenerated carriers.

Section 5.1 is a general description of the test chip. Descriptions of the circuits on the chip and procedures for measurement of the operating frequency of a transimpedance receiver circuit are presented with laboratory results in Section 5.2. The circuits, procedures and data for the response skew and leakage current tests are presented in Section 5.3 and Section 5.4, respectively.

5.1 General Chip Description

Figure 26 is a photograph of the CMOS optical clock receiver test chip. This chip contains three separate sets of test circuits. The 18 vertical strips of devices on the chip are individual optical receivers. The nine receivers in the upper half have the two-phase clock driver circuits included, while the circuits in the lower half do not. The receivers in the lower half are included to help align the optical input beam to the photodiode window for the leakage current test; the outputs of these receivers do not provide information regarding the data being measured for the test. The photodiodes, shown circled in the photograph, are near the lower parts of the receivers and are 20×20 microns in size. The area surrounding the photodiode is masked by an aluminum sheet 100×100 microns with a window to allow light to penetrate to the photodiode while protecting the surrounding silicon from incidental illumination. This mask is part of the top metal interconnect layer, and is electrically connected to the ground bus.

The leftmost vertical string of devices in the upper half of the layout comprises the test circuit for measuring the maximum clock rate. The eight similar vertical strips to the right of this circuit are receivers used in the skew measurement test, while the nine vertical blocks in the lower half of the layout contain the leakage current measurement test. This chip has been fabricated with the DARPA-sponsored MOSIS 3-micron CMOS process. The die size is 6900×6800 microns, using a standard 40 pin package. While a 1.5- or 2-micron fabrication process could have produced device sizes with response times and parameter tolerances more indicative of the state-of-the-art in silicon CMOS technology, the 3-micron process was chosen because it represented an established technology with reasonably good reliability. At the time of fabrication, the MOSIS 2-micron fabrication process did not have an established delivery and reliability reputation.

The MOSIS facility returned 13 copies of the test chip, each with limited performance reliability. The analog nature of the transimpedance amplifier makes this design very sensitive to deviations of the transistor threshold voltages greater

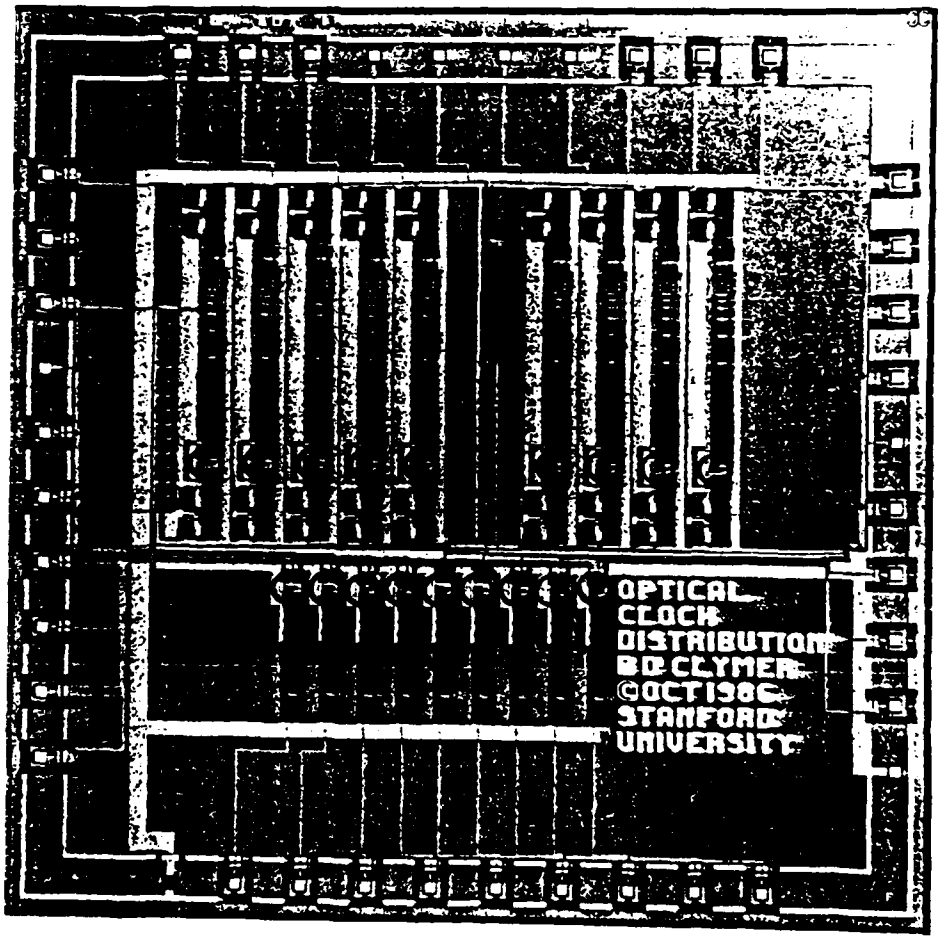

Figure 26: Optical clock distribution test chip.

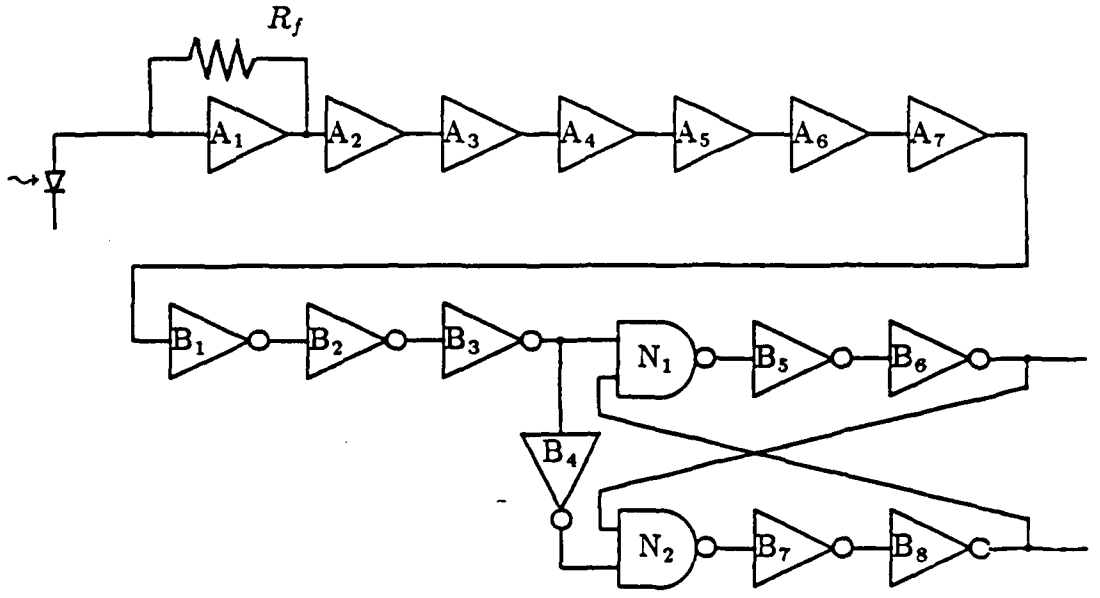

Figure 27: Transimpedance optical receiver and two-phase clock circuit for test chip.

than 100 mV from the nominal design values. Since the MOSIS CMOS process is primarily a digital circuit fabrication service, the threshold voltage range is larger than is generally acceptable for analog circuit implementation, and many of the receiver circuits on the test chip did not work as a result. The location of the receivers which failed varied from chip to chip, however, and enough working receivers were available over the 13 copies of the chip to perform measurements of the parameters under test.

The transimpedance receiver that has been presented in Chapter 3 is a principal element in each of the three test circuits on the chip presented in this chapter. The design presented in Chapter 3 has been scaled to accomodate a 3-micron fabrication process. A circuit diagram of the transimpedance receiver is shown in Figure 27. Devices labeled A_n are analog inverter stages, while those labeled B_n and N_n are digital logic gates. Table 6 contains device dimensions for the elements of Figure 27.

NAME	NMOS(μm)	PMOS(μm)	NAME	NMOS(μm)	PMOS(μm)
A_1	2070	720	B_1	24	24
A_2	2070	720	B_2	24	24
A_3	1035	360	B_3	63	123
A_4	552	192	B_4	63	123
A_5	276	96	B_5	246	492
A_6	138	48	B_6	246	492
A_7	69	24	B_7	246	492
			B_8	246	492
			N_2	63	123
			N_2	63	123

Table 6: Transistor sizes for amplifier and clock driver circuits.

The amplifier circuit in Figure 27 has a resistive feedback path to allow self-biasing and tracking of quiescent operating voltages for the chain. The photodiode for this particular chip design is a p^+diffusion into the n substrate fabricated in the same processes used to make sources and drains for the p-channel transistors in the standard MOSIS 3-micron p-well technology. The capacitive loading that is anticipated for distribution of each clock phase in the context of a practical VLSI system is modeled on the test chip by driving a large inverter stage with each clock output signal.

5.2 Operating Frequency Test

The output pads and pin connections for the dual inline package (DIP) in which the test chip is mounted are limited to frequencies lower than around 12 MHz. In order to measure on-chip clock frequencies higher than the bandwidth of the output pads on the test chip, a frequency divider circuit has been included. This circuit allows the indirect measurement of the on-chip clock signal frequency by observing a test signal with a frequency less than 1 MHz.

The test circuit for measuring the clock rate of the optical receiver is composed of an optical receiver and clock driver circuit described in the previous section and

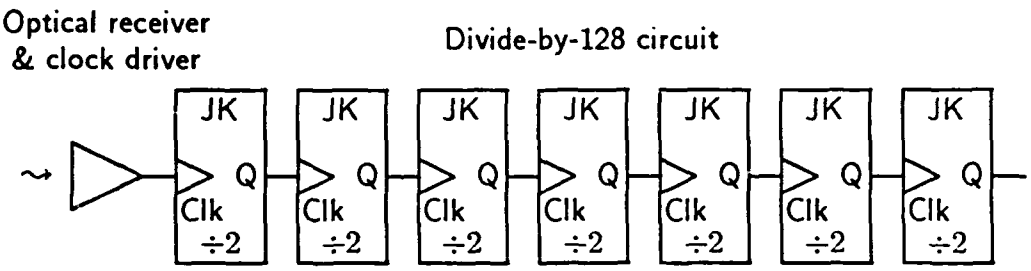

Figure 28: Frequency divider circuit for off-chip measurement of output frequency.

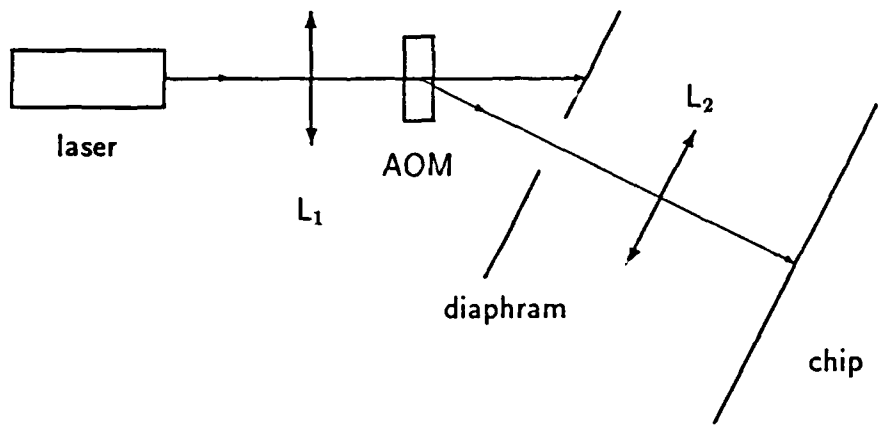

Figure 29: Optical layout for frequency measurement.

a divide-by-128 circuit. A block diagram of the circuit is shown in Figure 28. The output of one of the phases of the clock driver circuit is used as a clock input to the divide-by-128 circuit. The divide-by-128 circuit is a ripple counter circuit composed of successive master-slave JK flipflops that have been configured to toggle the output state for each clock cycle. Each stage divides the frequency of its clock input by a factor of 2[45].

A diagram of the optical layout used for the operating frequency measurement is shown in Figure 29. Since only a single spot of light is required for the clock input for the test, a hologram is not required. Instead, a single laser beam is focused to a spot on the chip surface. In the test configuration, a continuous wave

Ser. No.	Maximum Input Frequency (MHz)	Ser. No.	Maximum Input Frequency (Mhz)
1	16.8	7	18.0
3	14.7	9	13.2
4	18.5	11	16.0
5	12.2	12	14.1
6	15.2	13	12.0

Table 7: Maximum operating frequencies for sample of chips.

(CW) He-Ne laser beam is passed through an acousto-optic modulator (AOM). L_1 is a lens used to focus the laser beam to a narrow waist for interaction with the acousto-optic crystal, as recommended in the application notes for the device [46]. By spatially bandpass filtering the 1st order beam out of the acousto-optic modulator by means of a diaphram, the AOM can be used as a shutter. In this manner, the light beam incident on the chip surface is modulated on and off. The spot is imaged by lens L_2 onto the photodiode on the test chip. The frequency and waveform of the optical signal can be controlled by an electronic input signal to the acousto-optic modulator. For the tests in this section and Section 5.3, the input signal was a square wave of 50 per cent duty-cycle.

The operating frequency tests were performed by applying an optical square wave of 150 μW peak power to the photodiode corresponding to the transimpedance receiver which provides the input signal to the frequency divider circuit. The frequency of the optical signal, f_{in}, was increased until the low frequency output signal of the JK flip-flop circuit was no longer a TTL level signal with a frequency equal to $f_{in}/128$. The measured maximum operating frequencies for the copies of the chip tested are shown in Table 7. The mean value of the maximum operating frequency was 15.1 MHz, while the variance was 1.95 MHz^2. The receivers for the operating frequency test on serial numbers 2, 8 and 10 did not work. The total sample included measurements on 10 chips.

While the operating frequency measurements are based on a sample size too small to have statistical significance, the results give an indication of the typical operating frequency that is possible with a transimpedance receiver implemented

in 3-micron CMOS technology. The frequencies measured confirm that the trans-impedance receiver approach for optical clock distribution cannot give the high frequency performance required to represent an improvement over existing electronic systems using metal wire paths.

The measured results are also within a range of frequencies that might be expected for a 3-micron implementation scaled from the 2-micron simulations performed in Chapter 3. The 2-micron simulations showed operation up to 50 – 67 MHz for TTL-level output voltages. Scaling these frequencies down by a factor of 1.5 to reflect the larger transistors in the 3-micron test circuits, the expected maximum operating frequencies are 30 – 45 MHz. The measured values are lower than simple scaling predicts, but are within a factor of 2 or 3.

5.3 Receiver Skew Test

The receiver skew test is a measurement of the difference in response times of identical optical receiver and clock driver circuits that are fabricated on the same chip. Since chip fabrication causes variation of transistor channel dimensions and threshold voltages, the characteristics of identical circuits on a given chip vary slightly. This variation is generally observed to be more severe as the separation between the circuits increases. In this test, the clock driver output signals of eight identical circuits are compared in pairs. The skew between receiver outputs is measured by generating a pulse of width corresponding to the separation between the rising (or falling) pulse edges for the two receiver output signals. This skew pulse is created by using the two receiver outputs as inputs to an exclusive-OR circuit.

A diagram of the circuit used to determine the skew between the output signals of the clock driver pairs is shown in Figure 30. The four logic gates at the top of the figure perform an exclusive-OR operation on the output signals of two optical receiver and clock driver circuits. Since the exclusive-OR output is true when one and only one input is true, two pulses of width corresponding to the time

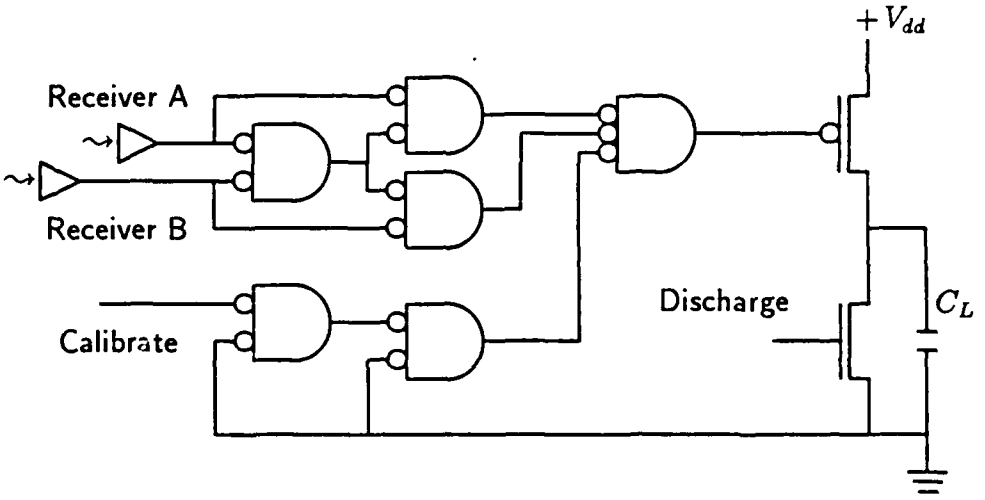

Figure 30: Receiver skew measurement circuit.

separation of the clock driver outputs are produced; the first pulse is formed by the separation of the leading edges of the two input signals, while the second pulse is formed by the separation of the trailing edges of the pulses. The exclusive-OR output pulses are negative-true logic levels, while the inputs are positive-true. The load capacitor, C_L, is charged by a PMOS pass transistor when the output of the exclusive-OR circuit is low. If the pulse is short enough, the capacitor does not charge entirely, and the final voltage across the capacitor is a measure of the charging time, and therefore a measure of the skew pulse width and the skew between the receiver output signals. A separate path is provided for calibration of the system, and a discharge input allows residual charge on the load capacitor to be removed to initialize the system for testing.

This system allows easy calibration of chip package and oscilloscope probe parasitics, because the charging characteristics of the load capacitor include the loading effects of the oscilloscope and probe. Since the exact RC value is not important, the added loading of the off-chip measurement equipment does not effect the accuracy of the measurement. The measurement is calibrated by observing the charging characteristic of the capacitor (including oscilloscope and probe) when a

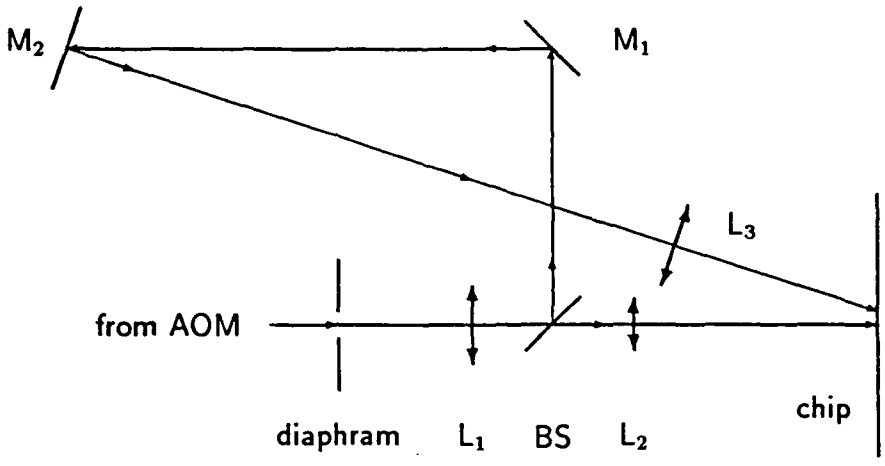

Figure 31: Optical layout for receiver response skew measurement.

step input is applied to the calibration input lead. The measured capacitor voltage during the test is compared to the point on the calibration charging characteristic, and the corresponding time to charge the load capacitance to that voltage is recorded. This charging time corresponds to the width of the skew pulse, and therefore to the skew between the receiver output signals. The resolution of the exclusive-OR gate is about four times the propagation delay for a typical logic gate.

The optical layout for the skew test is shown in Figure 31. This layout is similar to that used for the operating frequency test shown in Figure 29 in that an acousto-optic modulator is used as a shutter to modulate the CW laser input. The first order output of the AOM is first collimated by lens L_1, then split into two beams. One beam is extended in length by means of mirrors, and both beams are focused onto the chip surface by lenses L_2 and L_3. Each spot is focused onto the appropriate photodiode for the corresponding receiver and clock driver circuit in the pair being compared.

The added optical pathlength of one of the beams allows verification of the

Ser. No.	Receiver pair	Skew (ns)	Ser. No.	Receiver pair	Skew (ns)
1	0,4	10.0	3	0,4	15.0
3	1,5	20.0	3	2,6	15.0
3	3,7	20.0	4	0,4	20.0
4	1,5	15.0	5	2,6	8.0
5	3,7	8.0	7	2,6	12.0
9	0,4	17.5	9	1,5	10.0
9	3,7	10.0	13	1,5	15.0
			13	2,6	5.0

Table 8: Measured receiver response skew between transimpedance receiver and clock driver pairs.

skew measurement. The difference in pathlengths corresponds to an external delay of 5.5 ns for one of the optical input signals. This artificially extends the skew between the outputs of the two receivers by a known amount. The measured skew is actually the sum of the external skew due to the difference in optical pathlengths and the receiver skew due to fabrication process variation of transistor parameters. An accurate measurement of the actual receiver response skew can be performed by recording the skew time for the optical beams focused to the two photodiodes in one configuration, then reversing the configuration and observing the results. The actual receiver response skew is the average of the two measurements. The accuracy can be confirmed by verifying that the difference between the two measurements is twice the difference in the optical pathlengths, or 11 ns.

The measured data are shown in Table 8. For these measurements, the peak optical power of the spots received through the mask window was 120 μW each. The optical signal was a very low frequency square wave with 50 per cent duty-cycle. This allowed a stable continuous trace on the oscilloscope with a period between pulses which was long enough for the stored charge on the load capacitor to decay without need for external discharge to reset the test. The average skew between the receiver pairs sampled was 13.25 ns and the variance was 2.26 ns^2.

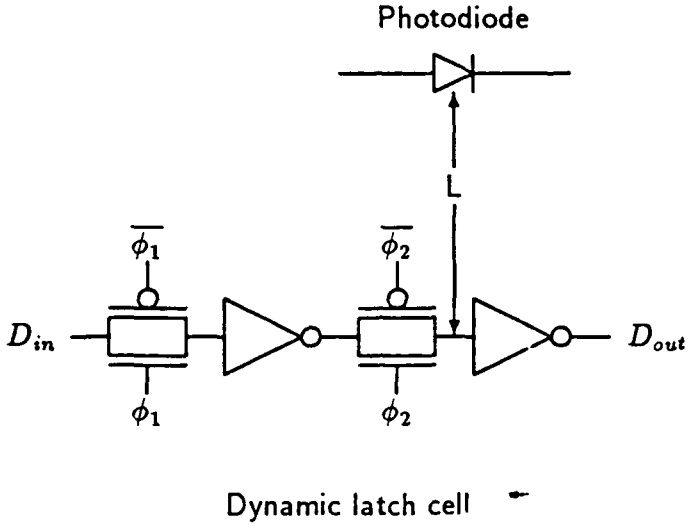

Figure 32: Test circuit to measure leakage current between photodiode and nearby storage cell.

The range of measured skew values was 5 – 20 ns. The skew data between the receiver pairs not listed in Table 8 were unavailable because of threshold failure of one or both of the receivers in the pair. The mean skew measured is about four times the expected propagation delay for a logic gate for the transistor sizes used in the test circuit.

5.4 Leakage Current Test

The leakage current test is composed of nine optical receivers, each with a dynamic latch cell positioned near the photodiode. A diagram of a typical test circuit for determining leakage currents is shown in Figure 32. The separation between the photodiode and the latch cell is varied from one receiver to the next. The receivers in this test have been designed with guard rings around the photodiodes and the associated latch cells. The photodiode is part of a transimpedance receiver as is the case in the other test circuits on the chip; however, the output of the optical

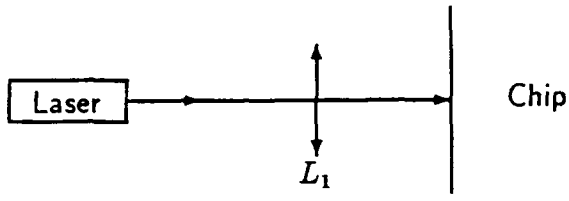

Figure 33: Optical layout for leakage current measurement.

receiver for this test is used only to aid in the alignment of the optical beam with the photodiode mask window.

The optical layout for the leakage current test is shown in Figure 33. Since no modulated signal is required for this test, the acousto-optic modulator is not used. The beam is focused directly onto the photodiode mask window. The beam can be blocked by an opaque screen to allow measurement of the dark storage time of the latch cell on the chip.

In this test, the storage time of each dynamic latch is observed as the photodiode is illuminated, and compared to the dark storage time to give an indication of the leakage currents caused by diffusion of carriers generated by deeply penetrating photons to transistors in the latch. The different spacings between latch cells and photodiodes allow quantification of the relationship between the leakage current strength and the distance carriers must diffuse from the substrate under the photodiode to reach the depletion layer of the latch cell pass transistor source region.

If the pass transistor is modeled as a perfect capacitor and the leakage currents are modeled as independent current sources, then the discharging time of the capacitor indicates the strength of the total current which transports charge away from the transistor source area. For normal operation, the pass transistor has a very small leakage current, and the charge remains on the source for a number of milliseconds before decaying. An additional leakage current due to carriers diffusing to the source area depletion region from the photodiode causes

Ser. No.	Separation (μm)			
	18	27	37.5	60
1	0.4	1.24	2.6	5.0
2	1.6	2.0	3.25	7.2
3	0.5	1.48	2.25	5.7
4	0.5	0.7	3.0	5.0
5	0.45	1.2	2.25	4.5
6	0.55	1.2	2.5	6.5
7	0.175	0.9	1.75	4.5
8	0.3	0.9	1.8	5.0
9	0.025	0.55	1.5	3.6
10	0.42	1.0	2.0	4.2
11	1.0	1.6	3.0	5.5
12	0.4	0.84	2.0	5.2
13	0.15	1.4	2.25	5.5

Table 9: Storage times in microseconds for latch cell with 2.2 mW optical power incident on photodiode at various separations to latch cell.

the charge to dissipate from the pass transistor at a much higher rate. This rate is proportional to the number of carriers present at the photodiode, and therefore is a linear function of the optical power incident on the photodiode. If the separation between the photodiode and nearby transistor depletion regions is large enough, many of the carriers which diffuse in the direction of these devices recombine before reaching the nearby transistor. Increasing the separation allows more carriers to recombine; therefore, the leakage current is reduced.

Table 9 lists the storage times in microseconds for the range of spacings between photodiodes and latch cells for the sample of test chips. A plot of these values is shown in Figure 34. These measurements were performed using an optical wavelength of 633 nm. The storage times would be considerably shorter for an incident optical wavelength 730 nm, since a much larger percentage of photons would penetrate into the substrate beyond the depletion layer of the photodiode and there would be therefore more carriers to diffuse to the nearby latch cell. The maximum available optical power, 2.2 mW, was used for the measurements in this

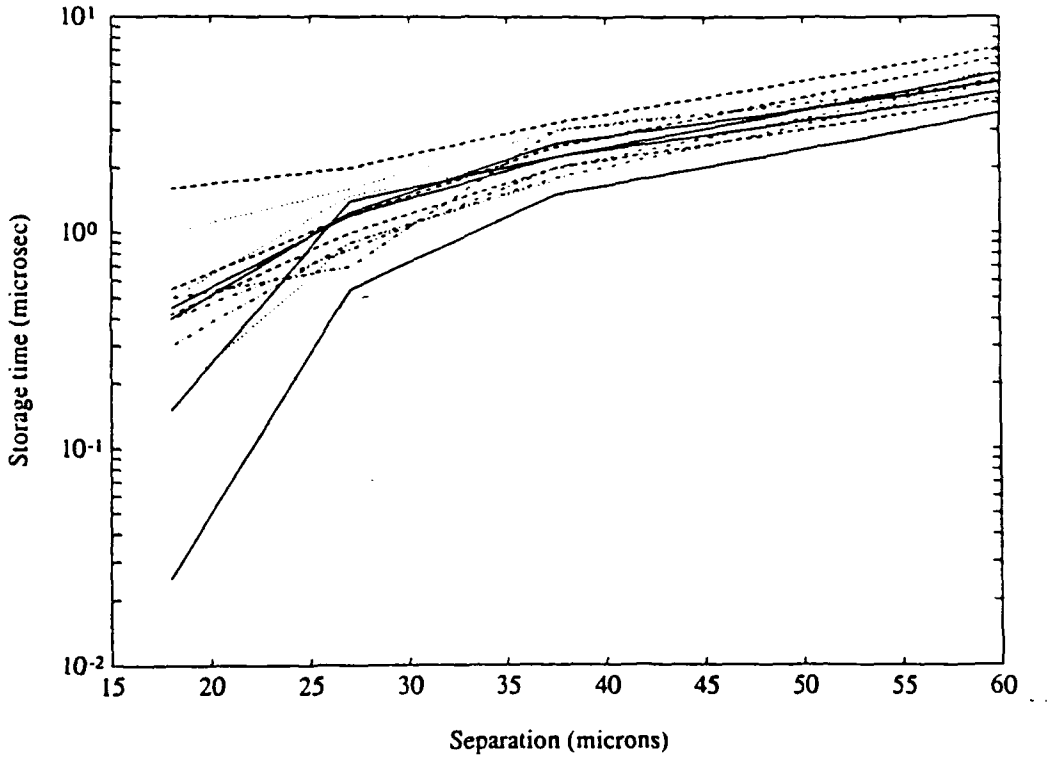

Figure 34: Storage times for latch cell as a function of separation between photodiode and latch cell.

test. Since the leakage current increases linearly with the incident optical power, the storage time is inversely proportional to the optical power, and the storage time for a lower level of input light can be interpolated by

$$t_2 = \frac{t_1 P_{opt,1}}{P_{opt,2}} \tag{85}$$

where t_1 is the known storage time corresponding to an incident optical power $P_{opt,1}$, $P_{opt,2}$ is the new power level, and t_2 is the new storage time. Equation 85 is valid as long as the dark leakage current through the pass transistor is small compared to the leakage current due to the illumination. Using the data from serial number 1 in Table 9 and the relationship in Equation 85, an incident optical power of 100 μW, has an expected storage time of around 9 μs for a separation of 18 microns, and a storage time of 0.1 ms for a separation of 60 microns. By limiting the incident optical power and requiring a minimum spacing design rule, the storage time for dynamic latch cells in the vicinity of a photodiode can be maintained to within system specification.

The test chip and data that have been presented in this chapter demonstrate some of the limits in the usability of a transimpedance optical receiver for clock distribution that were anticipated based on the initial investigation presented in Chapter 3. The layout area cost per receiver is very large, while the maximum operating frequency for a 3-micron implementation has been shown to be limited to a range of 10 – 20 MHz. While a 2-micron implementation would be expected to operated at a higher frequency, the improvement would not be greater than a factor of 2; therefore, a 2-micron design could be expected to have a maximum operating frequency less than 40 MHz. In addition, the response skew was found to be appreciable. The measured response skew range was 5 – 20 ns, with 20 ns representing about 6 times the average propagation delay for a logic gate on the chip. Clock distribution systems based on metal wiring can provide coverage for a VLSI chip with appreciably less skew, indicating that the transimpedance receiver for optical clock distribution cannot provide improved skew performance. These limits are not fundamental to the optical clock distribution approach, but rather to the type of optical receiver that has been implemented for the test chip.

A phase-locked loop receiver such as that presented in Chapter 4 allows improved performance with less layout area cost.

The leakage current data presented in this section represent a limitation of the optical clock distribution approach regardless of receiver type, when the constraints of implementing the photodiode as a vertical p-n junction with existing CMOS fabrication steps and bias voltages are observed. The problem is increased as the wavelength of illumination is increased, due to the larger absorption length of photons for longer wavelengths of light. Based on the data presented in this section, the storage time for dynamic latch cells can be preserved at a cost of increased separation between devices or a reduction in optical power. This also indicates that a phase-locked loop receiver is preferable since operation of this type of receiver requires much lower optical input power.

Chapter 6

Conclusions

A design and analysis of an optical clock distribution system have been presented. A summary of contributions is given in Section 6.1, followed by suggestions for future work in Section 6.2.

6.1 Summary of Contributions

This investigation has involved the design of an optical clock distribution system for silicon VLSI computer chips. A general description of the clock distribution system has been presented in Chapter 2 along with the goals and constraints for the system design and the engineering trade-offs that are inherent in them. This system uses a free-space three-dimensional approach in which the optical clock signal is mapped from an off-chip light source via a holographic or lenslet array element to several photodetectors integrated on the silicon chip surface. In the analysis presented in this thesis, the effect of crosstalk from secondary reflections is not considered.

In Chapter 3, the first of two receiver designs has been presented. The transimpedance receiver presented has been adapted from a well-established optical communication approach. It has been shown by circuit simulation experiments and

noise jitter analysis, however, that an implementation of this receiver in CMOS for integrated detection of optical clock signals does not offer a substantial improvement in performance over existing metal wire clock distribution methods. The maximum operating frequency for a 2-micron implementation of the transimpedance receiver and clock driver is approximately 50 MHz, and is limited by the response skew between receivers implemented in different locations on a given chip.

The poor performance of the transimpedance receiver has led to the development of the phase-locked loop receiver presented in Chapter 4. Circuit simulations presented show that this receiver can operate at frequencies of 100 – 200 MHz with picosecond steady-state response skew. In addition, the noise analysis presented shows that the phase-locked loop can operate with sub-picosecond signal jitter. The comparison of optical power requirements for the two receiver examples presented shows that the phase-locked loop receiver can meet system timing specifications with much lower incident optical power levels than the transimpedance receiver.

A CMOS chip to measure three parameters of an integrated optical receiver and photodiode have been presented in Chapter 5. This test chip includes tests to quantify the maximum operating frequency and response skew for a typical transimpedance receiver with 3 micron channel length transistors, as well as measure the leakage current between an integrated photodiode and nearby dynamic storage devices. The maximum operating frequency of the transimpedance receiver has been measured to be approximately 15 MHz for the 3-micron implementation, while the skew between output signals of paired receivers ranges from 5 – 20 ns. The leakage current between the carriers generated in the substrate under the photodiode can be reduced by limiting the incident optical power or requiring a minimum separation between the photodiode and sensitive devices. A storage time for a nearby dynamic latch cell of at least 0.1 ms can be maintained with a minimum spacing of 60 microns and a maximum optical input power of 100 μW.

6.2 Suggestions for Future Work

The work presented here represents an initial effort in demonstrating the feasibility of optical clock distribution for VLSI chips. The phase-locked loop receiver presented in Chapter 4 should be implemented in a test chip and the operating parameters of the receiver such as the conversion gains, the range of input frequencies to which to loop will lock, and the difference in steady-state phase errors between identical receiver copies should be measured.

In addition to measuring the phase-locked loop receiver parameters, a test configuration in which the clock is distributed optically to several receiver sites on a chip should be implemented. Such a configuration would require the development of an optical mapping element, and would give an indication of the production and alignment problems inherent in the incorporation of a hologram or lenslet array into the optical distribution system. The electronic logic circuits on the chip receiving the optical clock signal should be an existing CMOS system, adapted to include phase-locked loop optical receivers and local polysilicon clock distribution within functional cells.

A further investigation of the leakage current and response time for the photodiode using near infrared illumination should be performed. Since a realistic package configuration for the optical clock distribution system would probably involve a semiconductor optical source, the detection parameters for illumination at wavelengths available with laser diodes should be measured, so that the feasibility of using coherent light and a holographic mapping element can be determined.

Finally, an investigation of the effects of secondary reflections of the light beam should be performed. Since the analysis presented here neglects these effects, this assumption should be verified or the effects should be characterized so that a more detailed analysis can be performed.

Bibliography

[1] J. W. Goodman. Optical interconnects in microelectronics. *Proc. SPIE Conf. on Optical Computing*, 456:72–85, 1984.

[2] H. Tajima, Y. Okada, and K. Tamura. A high speed optical common bus for a multiprocessor system. *Trans. Inst. Electron. and Commun. Eng. Japan*, E66:47–48, 1983.

[3] R. K. Kostuk and J. W. Goodman. Optical imaging applied to microelectronic chip-to-chip interconnections. *Appl. Opt.*, 24(17):2851–2858, 1985.

[4] J. W. Goodman, F. Leonberger, S. Y. Kung, and R. Athale. Optical interconnections for VLSI systems. *Proc. IEEE*, 72(7):850–866, 1984.

[5] R. A. Milano. GaAs integrated optoelectronic circuits. In *Winter '85 Topical Meeting on Optical Computing, Technical Digest*, page MB1, Optical Society of America, 1985.

[6] F. Anceau. A synchronous approach for clocking VLSI systems. *IEEE J. Solid-State Circuits*, SC-17(1):51–56, 1982.

[7] S. Dhar, M. A. Franklin, and D. F. Wann. Reduction of clock delays in VLSI structures. In *Proc. IEEE 1984 Conf. on Computer Design: VLSI in Computers*, pages 778–783, 1984.

[8] R. W. Keyes. Communication in computation. *Intl. J. Theor. Phys.*, 21(3,4):263–273, 1982.

[9] K. C. Saraswat and F. Mohammadi. Effect of scaling of interconnections on the time delay of VLSI circuits. *IEEE Trans. on Electron Dev.*, ED-29(4):645–650, 1982.

[10] Carver Mead and Lynn Conway. *Introduction to VLSI Systems*, pages 230–231. Addison-Wesley, 1980.

[11] Carver Mead and Lynn Conway. op. cit. page 51.

[12] Carver Mead and Lynn Conway. op. cit. page 164.

[13] Carver Mead and Lynn Conway. op. cit. page 14.

[14] Ronald N. Bracewell. *The Fourier Transform and Its Applications*, page 345. McGraw-Hill, 2nd edition, 1978.

[15] D. F. Wann and M. A. Franklin. Asynchronous and clocked control structures for VLSI based interconnection networks. *IEEE Trans. Computers*, SC-19(3):284–293, 1983.

[16] Carver Mead and Lynn Conway. op. cit. pages 242–246.

[17] H.-H. Witte. Optical LAN activities in Europe. *J. Lightwave Tech.*, LT-3:432–437, 1985.

[18] E. G. Rawson. The Fibernet II ethernet-compatible fiber-optic LAN. *J. Lightwave Tech.*, LT-3:496–501, 1985.

[19] Davis H. Hartman. Digital high speed interconnects: a study of the optical alternative. *Opt. Eng.*, 25(10):1086–1102, 1986.

[20] L. A. Bergman, W. H. Wu, A. R. Johnston, R. Nixon, S. C. Esener, C. C. Guest, P. Yu, T. J. Drabik, M. Feldman, and S. H. Lee. Holographic optical interconnects for VLSI. *Opt. Eng.*, 25(10):1109–1118, 1986.

[21] J. A. Fried. Optical I/O for high speed CMOS systems. *Opt. Eng.*, 25(10):1132–1141, 1986.

[22] Raymond K. Kostuk. *Multiple Grating Reflection Volume Holograms with Application to Optical Interconnects.* PhD thesis, Stanford University, Stanford, California, 1986.

[23] P. R. Haugen, S. Rychnovsky, A. Husain, and L. D. Hutcheson. Optical interconnects for high speed computing. *Opt. Eng.*, 25(10):1076–1085, 1986.

[24] M. K. Kilcoyne, S. Beccue, K. D. Pedrotti, R. Asatourian, and R. Anderson. Optoelectronic integrated circuits for high speed signal processing. *Opt. Eng.*, 25(10):1119–1126, 1986.

[25] B. D. Clymer and J. W. Goodman. Optical clock distribution to silicon chips. *Opt. Eng.*, 25(10):1103–1108, 1986.

[26] A. A. Abidi. Gigahertz transresistance amplifiers in fine line NMOS. *J. Solid-State Circuits*, SC-19(6):986–994, 1984.

[27] S. K. Enam and A. A. Abidi. Decision and clock recovery circuits for gigahertz optical fiber receivers in silicon NMOS. *J. Lightwave Tech.*, LT-5(3):367–372, 1987.

[28] S. M. Sze. *Physics of Semiconductor Devices*, pages 749–752. John Wiley & Sons, 2nd edition, 1981.

[29] T. L. Kamins and G. T. Fong. Photosensing arrays with improved spatial resolution. *IEEE Trans. Electron Dev.*, ED-25(2):154–159, 1978.

[30] Carver Mead and Lynn Conway. op. cit. pages 232–233.

[31] P. R. Gray and R. G. Meyer. *Analysis and Design of Analog Integrated Circuits*, page 719. John Wiley & Sons, 2nd edition, 1984.

[32] P. R. Gray and R. G. Meyer. op. cit. page 749.

[33] Carver Mead and Lynn Conway. op. cit. Chapter 7.

[34] H. Taub and D. L. Schilling. *Principles of Communication Systems*, pages 298–300. McGraw-Hill, 1971.

[35] Floyd M. Gardner. *Phaselock Techniques*. John Wiley & Sons, 2nd edition, 1979.

[36] Ronald E. Best. *Phase-Locked Loops*. McGraw-Hill, 1984.

[37] Floyd M. Gardner. op. cit. page 11.

[38] Floyd M. Gardner. op. cit. pages 30–31.

[39] Ronald E. Best. op. cit. pages 36–41.

[40] Floyd M. Gardner. op. cit. page 55.

[41] Floyd M. Gardner. op. cit. page 43–44.

[42] S. M. Sze. op. cit. page 111.

[43] S. M. Sze. op. cit. page 753.

[44] S. M. Sze. op. cit. page 758.

[45] Gerald E. Williams. *Digital Technology*, pages 203–204. Science Research Associates, 1977.

[46] *Application Note: Acousto-Optic Modulators Model 3000 Series*. Crystal Technology, Inc., Palo Alto, California.

APPENDIX B

AN IMPLEMENTATION OF A PHASE-LOCKED LOOP APPROACH TO THE PROBLEM OF OPTICAL CLOCK DISTRIBUTION FOR SILICON VLSI

A THESIS
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
ENGINEER

By
Dick Philip Welch
December 1988

Approved for the Department:

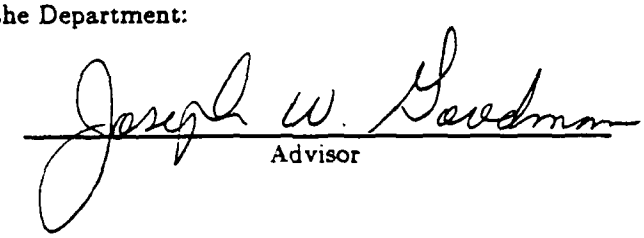

Advisor

Approved for the University Committee on Graduate Studies:

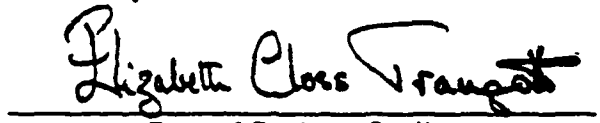

Dean of Graduate Studies

Contents

List of Figures

Acknowledgements

Like any project, this one required the help and advice of many people to whom I owe a debt of thanks. First of all is my advisor, Professor Joseph Goodman. He gave me the freedom to learn and do on my own, but was there when I needed advice or equipment. He showed me that academic success and teaching ability are not necessarily mutually exclusive. Finally, he was not ashamed to admit ignorance of a subject, and yet could still point me to a person possessing the knowledge I required. For all of this, I give my thanks.

Second, I wish to thank James Burr, Ernst Kimler, and the rest of the SETI group at Stanford. Without the use of their facilities, the layout and fabrication of the test chip would not have been possible. Also, their design suggestions and computer help have been invaluable.

Third, Ms. Karin Sperley deserves many thanks for graciously agreeing to complete testing of the OCD2 chip after other commitments forced me to leave Stanford before I could finish testing it.

I also wish to express my thanks to Professor Bruce Wooley for the discussions of some of the designs considered and implemented in this project.

To Professor David Bloom, two special votes of thanks are due. The first of these is for the loan of a signal generator with which to modulate the diode laser's optical clock signal. The second and by far more important thanks is for displaying a genuine concern for the welfare of students. Whereas almost every professor shows concern for his or her own students, Professor Bloom

has gone out of his way to help me on several occasions; and discussions with other students reveal that I am not unique in this regard. This rare concern for students has earned Professor Bloom both my sincere thanks and my utmost respect.

To Dr. Kenneth Fesler, Robert McRuer, and Larry MacAdams: I wish to express my gratitude for the help in designing the optical setup necessary for the testing of the OCD2 chip. And to the rest of Profs. Goodman's and Hesselink's groups, a thanks is due for the loan of the optical equipment necessary to test this chip.

Next, I thank my father, for teaching me to be the harshest critic that inspects my own work; and I thank my mother for teaching me to be lenient when examining another's work.

And for various reasons of technical, practical, and/or psychological support, I wish to thank the following: Ms. Sherri Hake, the crew at MOSIS, the Thug Brothers, Bruce Springsteen and the E-Street Band, Glen Dudevoir, the Spikeheads, Meemer Bleamer, all of SWUNET, the Nylons, Van Morrison, the Poman, the Stanford Self-Defense Club, the St. Louis Cardinals, Professor Macovski's group, the Chicago Blackhawks, and (last but not least) the Blowfish Co-ed Inner Tube Water Polo Team.

And now, a few words from our sponsors:

This material is based upon work supported under a National Science Foundation Graduate Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

This work is supported under a grant from the Army Research Office.

Dick P. Welch
Stanford University
August, 1988

Chapter 1

Introduction

In conventional VLSI[1] design, the data path contains a combination of synchronous and asynchronous logic. Typically, combinational logic is separated by latches (registers) which hold intermediate results; and each latch is activated by a global clock signal.[2] Therefore, the clock signal is one of the most heavily loaded and most widely distributed signals on the chip surface. With current device sizes and geometries, the delay associated with propagation down a length of wire dominates the delay associated with switching devices [6, pp. 200–231]. Hence, the delay from when the clock actually switches until when the gate of interest switches is dominated by effects of the wire path connecting the gate to the clock driver. This delay is proportional to the length of the wire path and will therefore be different for gates located at different distances from the clock driver. The difference in "arrival times" of the clock signals at the various gates is referred to as clock skew, and the maximum clock skew is one limit on the operating frequency of the clock.

In a free-space optically distributed clocking scheme, the clock signal is gen-

[1] Very Large Scale Integration

[2] Actually, most clocking schemes feature multi-phased clocks. Therefore, strictly speaking, the latches are driven by phases of the clock.

erated (using an off-chip laser or LED) and mapped optically onto various sites on the chip's surface. Here, receivers convert the optical clock into an electrical clock signal and distribute it locally on the silicon surface. Since the speed of light is vastly greater than the speed of chip signal propagation, the optical portion of this system is essentially skew-free. With the chip divided into regions which are serviced by local clock receivers, the maximum distance a signal must electrically propagate is reduced. Hence, the maximum difference in clock signal arrival times is reduced and the system should be capable of operation at higher clock frequencies, assuming that the skew introduced (between local clocks) by the optical receivers is not excessive.

As detailed by Clymer [2], optically distributed clocking may indeed be a feasible alternative to electrically distributed clocking in regard to VLSI circuits fabricated on silicon. To be feasible for such implementation, an optically distributed clocking system should be able to match or exceed global system clock rates attainable using conventional electrically distributed clocking. Furthermore, to minimize the increased cost of such an optical clock scheme, it is desirable to be able to fabricate the necessary receiving circuitry using a standard silicon technology.[3] In addition, an optical clock scheme should avoid necessitating the use of regularly spaced optical receiver circuits, since such a requirement conflicts with standard VLSI design practice of laying out regular functional blocks and then placing these blocks in locations which attempt to optimize system performance but which are not necessarily arranged in a regular grid. Finally, the silicon area occupied by the optical clock receivers should be small enough that useful circuits can still be fabricated on the remaining silicon surface.

Clymer has made two attempts to design optical clock systems which meet the above criterion. The first design used an established transimpedance receiver

[3]Note that this precludes the use of p-i-n photodiodes, hybrid photodetector/VLSI packages, and/or special implants to modify the photosensitivity of the silicon surface.

to detect the optical signal and then used a series of analog stages to amplify the signal to TTL signal levels. A series of test chips were fabricated in the DARPA-sponsored MOSIS 3μm n-well CMOS technology. The test chips had an average maximum operating frequency of 15MHz and an average receiver skew of 13.25ns. Unfortunately, the transimpedance receiver circuits required a large amount of silicon surface area and did not show marked improvement over existing electrically distributed clocking approaches, and hence did not seem to represent a viable alternative to electrically distributed clocking.

Clymer's second design utilizes a phase-locked loop approach to the problem of receiving an optically distributed signal. In this approach, local oscillators located in various places on the chip surface are used to generate the clock signal; and the optical signal is used to synchronize these local oscillators' clocks via phase-locking. (See section 2.1 for details.) Since the optical signal is only used to synchronize the local oscillators, it doesn't need as much amplification as it would to drive the circuitry associated with a transimpedance approach. With less amplification, there is less of a chance of introduced skew (due to cross-chip process variations) between local clocks. Hence, the phase-locked loop approach should result in less skew and should therefore allow a higher operating frequency. Clymer's simulations indicate that, when fabricated in 2μm CMOS, such a design should be capable of operation with a maximum frequency in the range of 100–200MHz. It is the subject of this thesis to report on an attempted implementation of a modification of Clymer's design using the DARPA sponsored MOSIS 2μm n-well CMOS technology.

In summary, a phase-locked loop optical receiver has been designed subject to the criteria of using standard silicon fabrication technology, allowing random clock receiver locations, and of minimizing the use of silicon surface. This design is a modification of Clymer's proposed phase-locked loop design; such modifications being necessary to allow implementation in the chosen technology. Results of simulations of this circuit are provided, and test chips implementing

this design are fabricated. Finally, ideas for future work are suggested.

Chapter 2

Design of Test Chip

To test the phase locked-loop design, it was not only necessary to fabricate th local Optical Receiver/Clock Amplifier (ORCA) cells, but it was also necessary to provide diagnostic circuitry on the chip. This diagnostic circuitry is comprised of skew detection units,[1] frequency division units,[2] and functional test units.[3]

The Optical Clock Distribution II (OCD2) chip as submitted to MOSIS contains 9 ORCA units, 12 skewdetector circuit blocks, 1 multiplexer/frequency divider unit, and 4 shift register/test units. The 9 ORCA units form a 3 by 3 array; and between each two horizontally or vertically adjacent ORCA's is a skewdetector block. The other circuitry is located on the array's periphery. (See figure 2.1.)

Each circuit block will now be described in detail.

[1] The skew detectors are designed to check for skew between two adjacent ORCA units by comparing corresponding non-overlapping clock phases.

[2] The frequency dividers are provided to permit measurement of the ~100MHz clock frequency using output pins with a maximum operating frequency of approximately 12-20MHz.

[3] The functional test units are designed to insure that the clocks are capable of clocking useful computational circuitry, as well as to place a "load" on selected clocks.

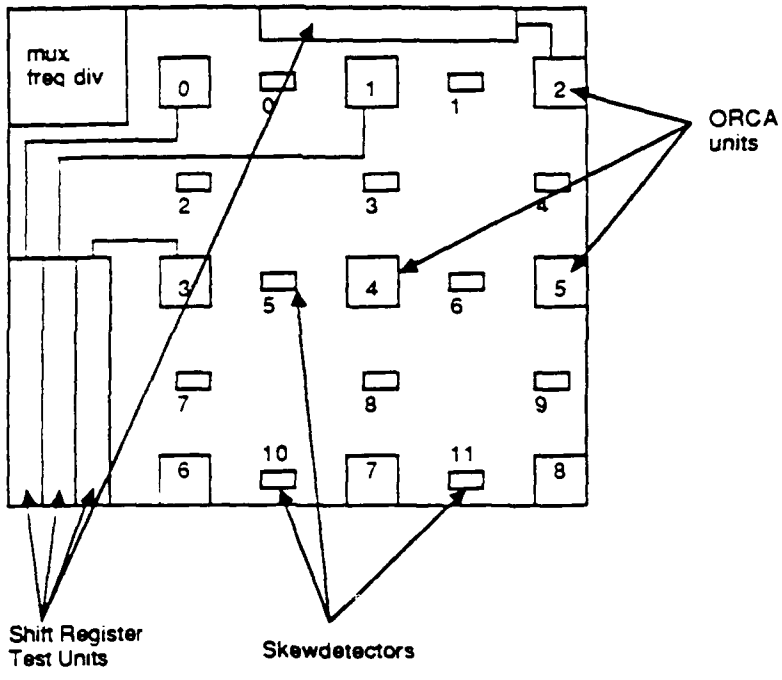

Figure 2.1: Block diagram layout of OCD2 chip.

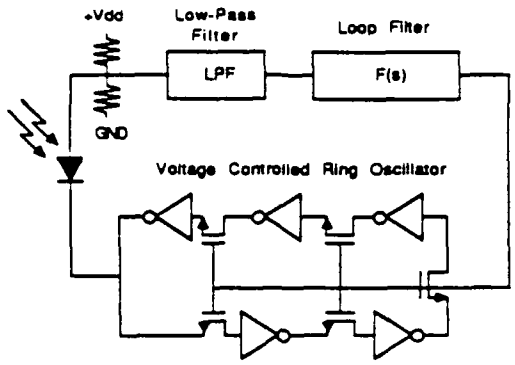

Figure 2.2: Block diagram of Clymer's phase-locked loop design

2.1 The Optical Receiver/Clock Amplifier cells

Each Optical Receiver/Clock Amplifier unit incorporates a phase-locked loop approach to the problem of receiving an optically distributed clock and distributing it locally on a silicon chip. The phase-locked loop used is a hybrid approach in which analog circuits are used for phase detection and filtering, while a digital circuit is used for the actual oscillator. A block diagram based on Clymer's work [2, pp. 44] is presented. (See figure 2.2.) However, due to restrictions imposed by the standard technology, several modifications to Clymer's design were necessary. These modifications will be noted as each subcircuit of the ORCA is now detailed.

2.1.1 The Voltage Controlled Oscillator

The VCO is basically a digital ring oscillator comprised of five inverters with one transmission gate in the loop. (See figure 2.3.) The transmission gate has its NMOS gate fed with a control signal CONTR, and the PMOS gate is driven by a signal CONTR$\sim$ generated by passing CONTR through a minimum size

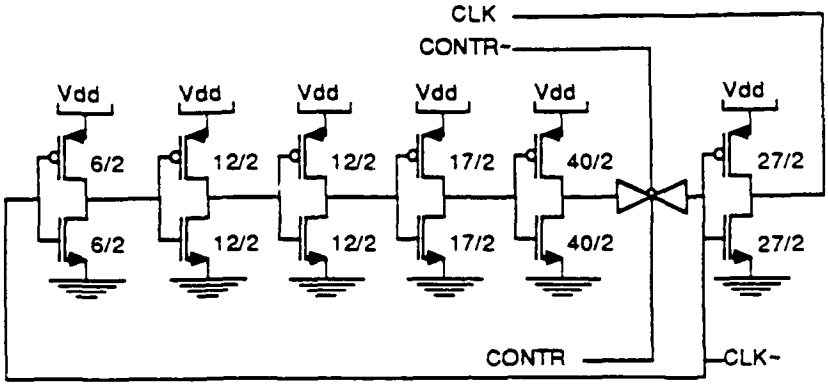

Figure 2.3: Basic digital voltage controlled oscillator

inverting amplifier. Hence, the signal CONTR acts as voltage control on the equivalent resistance of the transmission gate, which in turn acts as a control on the ring oscillator's frequency. SPICE simulations indicate that this design can provide an oscillator frequency in excess of 200MHz while insuring that the oscillator will not reach a stable state under worst-case conditions of input voltage and/or fabrication variations.[4]

The oscillator actually implemented has had its frequency adjusted[5] to be approximately 100MHz when the control voltage CONTR is approximately 4.2 volts.[6]

Note that this VCO design is slightly modified from Clymer's [2, page 44] in that four of the pass transistors have been removed and the fifth has been changed to a transmission gate. By changing the pass transistor to a transmis-

[4]Note that there must be an odd number of inverters in the ring oscillator to prevent it from settling to a steady-state. Furthermore, simulations indicate that the feedback of a ring oscillator with fewer than five inverters could result in its reaching a stable state (such as 2.5V) rather than the oscillation (between 0 and 5 V) which is desired.

[5]Adjustment is done via adjustment of the lengths and widths of the transistors in the transmission gate to affect its equivalent resistance.

[6]This is due to requirements imposed by the photodetector/amplifier block and will be explained later (on page 19).

sion gate, we necessitate the generation of a second control signal CONTR~.[7] However, we gain a significant speed advantage in the ring oscillator since the pass transistor suffers the well-known "threshold drop" problem, while the transmission gate can pass signals over the entire voltage range. Also, a pass transistor would require that CONTR always be above approximately 3.5 volts in order to keep the five inverters of the ring from becoming decoupled; but by using a transmission gate, the range of voltages over which CONTR can vary increases slightly.

Furthermore, removing four of the pass transistors increases the maximum operating frequency of the oscillator at the expense of decreasing the range of frequencies over which the loop can be operated. However, as we shall see, the photodetection circuitry already imposes a limit on the range of frequencies over which the oscillator can operate. Thus, we can increase the maximum operating frequency of the circuit without significant penalty (in terms of operating frequency range.)

By modifying Clymer's design as described above and by ratioing the inverter dimensions to provide increased drive current to the transmission gate while presenting it with a near minimal load,[8] the maximum frequency of the ring can be increased from approximately 67MHz[9] to over 200MHz, according to SPICE simulations.

2.1.2 The Loop Filter and Low Pass Filter

As noted by Clymer[2, pp. 44–53] and detailed in Gardner [3], the use of a simple lag-lead filter is well understood in phase-locked loop theory. Furthermore, it has reasonable layout requirements and a small steady-state phase error. Hence, Clymer's choice of a loop filter design is used here. As for the low pass filter,

[7] CONTR~ should be approximately V_{dd}–CONTR.
[8] Minimal in the sense of parasitic capacitance which must charge through the T-gate.
[9] Simulation of Clymer's design using the SPICE models indicated in Appendix B.

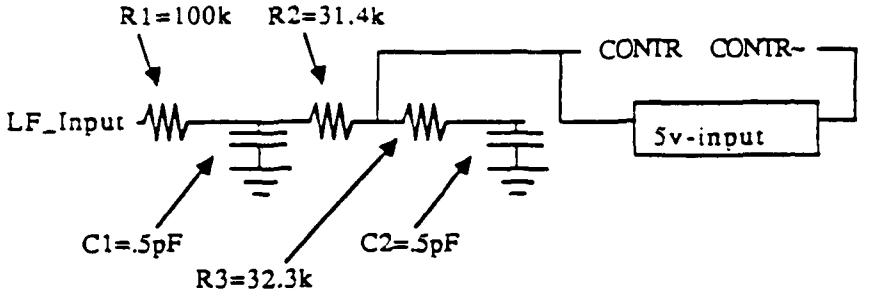

Figure 2.4: "Ideal" low-pass/lag-lead filter combination.

we are subject to three basic criteria: First, we need to filter out the 100MHz signal from the phase detector's comparison, for residual signal will result is unwanted variation on the control voltage of the VCO causing jitter in the VCO's output frequency. Second, we do not want to unduly increase the lock-in time[10] of the phase-locked loop by making the RC time constant of the low pass filter too large. Finally, we must be able to implement the filter using standard technology, which limits the range of component values which we may use.

In answer to these criteria, a simple RC low pass filter has been designed with its breakpoint set at τ = 20MHz. This provides adequate filtering of the 100MHz signal, and it can be implemented using a 100kΩ resistor and a 0.5pF capacitor. A diagram of the "ideal"[11] cascaded low pass/lag-lead filter is provided in figure 2.4, and a frequency response plot for this filter is shown in figure 2.5.

However, one must remember that we are working in a technology which does not easily provide capacitors and resistors—instead we use transistors to approximate the characteristics of these devices. (See figure 2.6 for transistor implementation of circuit.) In doing so, the linear approximation (of modeling

[10] Time from when the circuit is switched on until the VCO "locks" to the optical signal.

[11] "Ideal" is used in the sense that the circuit shown assumes ideal resistors and capacitors, rather than the transistors which must be used to implement them.

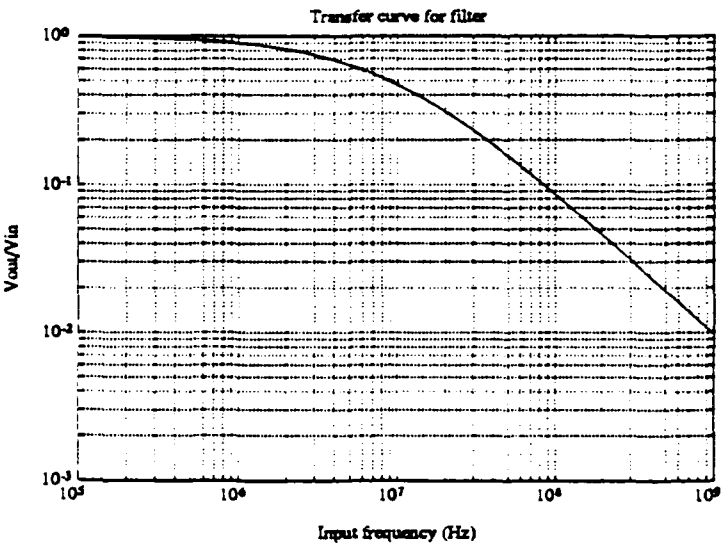

Figure 2.5: Frequency response of ideal filter network.

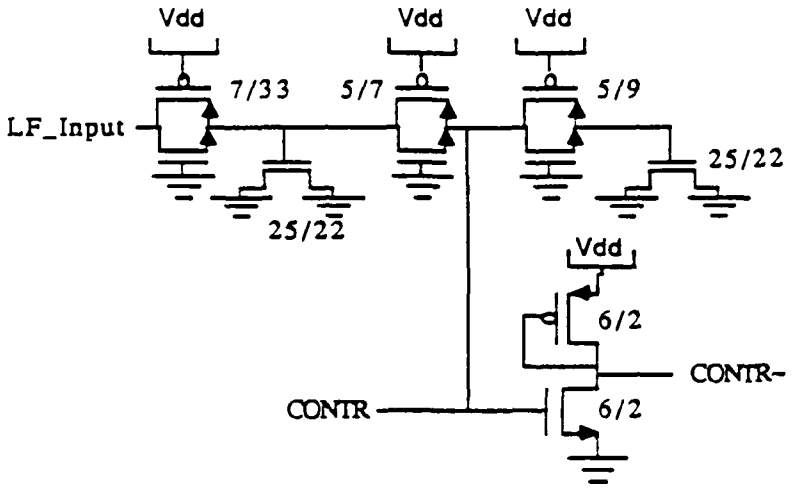

Figure 2.6: Low-pass/lag-lead filter combination as actually implemented on the OCD2 chip.

these transistors as resistors and capacitors) is reasonably valid over only a limited range of input voltages to the filter. Fortunately, we intend to operate the filter in such a way that the output voltage (CONTR) is approximately constant at 4.2 volts.[12] If we assume that the output is at a steady state value of 4.2 volts, the "capacitor" C_2 has a constant gate bias of approximately 4.2 volts, which indicates that it stays in the "linear" regime of operation and therefore has a constant capacitance [11, 4]. Also, as long as the "capacitor" C_3 has a gate voltage above its threshold voltage (approximately 0.8 volts), it too is in the linear regime of operation. Therefore, it is valid to model the transistors used to implement C_2 and C_3 as linear capacitors.

As for resistances, R_3 has a very small V_{DS} and a $\mid V_{GS} \mid$ of 4.2 volts.[13] Thus,

[12] This neglects the "power on" charging of the capacitors of the filter. When the chip is first powered on, the capacitors will be non-linear, but once they reache a steady state, the approximations mentioned above become valid.

[13] This is for the PMOS half of the resistor. The NMOS half has a gate-source voltage of about .8 volts, which means it is just barely "on" and has a high resistance. Hence the parallel

the "resistor" is in the linear regime, and has a nearly linear resistance over the small range of V_{DS} it experiences. Over the voltage range of interest, therefore, R_3 can be reasonably modeled as a linear resistor. R_2 is under conditions which are similar to those of R_3, and can reasonably be considered a linear resistor for identical reasons. R_1 is a more complicated case. Both the source and the drain of this resistor experience variations. However, the variations on the LF_Input side of the resistor are generally larger than those at the other side. Hence, we shall assume the voltage on the capacitor/resistor junction side of R_1 is constant at 4.2 volts. Even with this assumption, however, the voltage V_{DS} has a worst-case range of approximately (-3.2V,+1.0V). Over this range, R_1 exhibits non-linear characteristics. If one assumes that the loop locks to a value where the input voltage range is smaller, one can make linearization arguments, but this assumption is hardly valid in all cases. Furthermore, the voltage range is lessened by the loading of the network,[14] and this would tend to make the resistance characteristics of R_1 more nearly linear. However, R_1 is still basically non-linear in characteristic—we merely model it as linear to get an estimate of the filter's performance and then check its actual performace at the frequencies of interest using SPICE simulation.

Another factor to consider when designing the filter is that the SPICE models currently available from MOSIS appear to be off by a factor of two when modeling resistance and capacitance, according to Burr [1]. Hence, we must restrict the range of voltages entering the low pass filter to fall withing a range in which the linear approximation is reasonably valid and we must be careful that a "factor of two" error in capacitance/resistance values will not cause the circuit to malfunction. The two worst-case[15] transfer curves for the ideal[16] fil-

PMOS transistor's resistance dominates.

[14]The (-3.2V,+1.0V) range comes from SPICE simulations from the unloaded output of the NAND gate comparator described in the next section. It is reasonable to assume that this voltage range would be decreased by the loading of the NAND.

[15]subject to the constraint of keeping RC constant, since the charging time estimates provided by MOSIS appear to be accurate.

[16]i.e. comprised of capacitors and resistors rather than transistors.

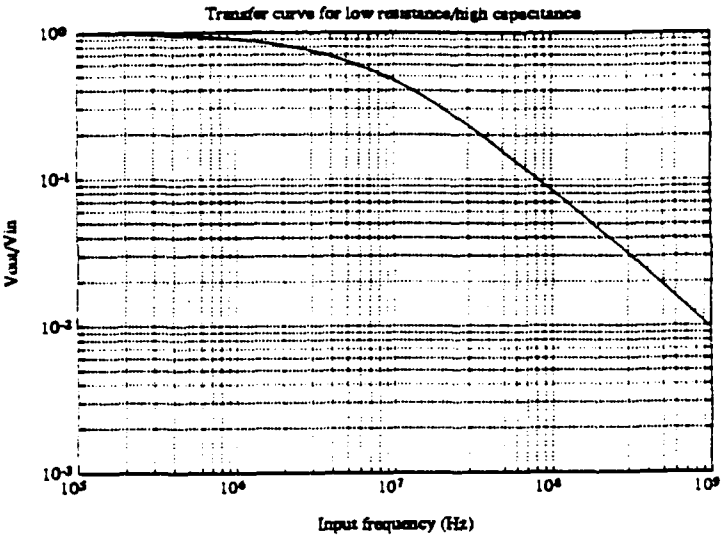

Figure 2.7: Transfer curve for low resistance/high capacitance values.

ter are shown in figures 2.7 and 2.8; and they appear identical to the nominal ideal filter. As the linear approximation of the current design appears adequate, SPICE simulations of the design's response to sinusoidal inputs of 100MHz and 12.8MHz[17] were performed and the results are presented in figures 2.9 and 2.10.

In the SPICE simulation of the 12.8MHz signal, we can see that it is attenuated somewhat by the filter. However, to keep the 100MHZ signal well filtered, we cannot easily reduce the attenuation of the 12.8MHz signal. Thus, there is some unavoidable increase in lock-in time of the loop (over that indicated by Clymer's simulations.) However, once the clock is locked to the optical signal,

[17] Clymer [2, page 52] calculated a "pull in" range of 12.8MHz for his original design. The 12.8MHz signal plotted here is not severely attenuated by the implemented low-pass/lag-lead filter combination. Thus, the "pull in" time of the loop is not excessively increased by our choice of low pass filter and by non-linear effects of the transistors used to implement said filter combination.

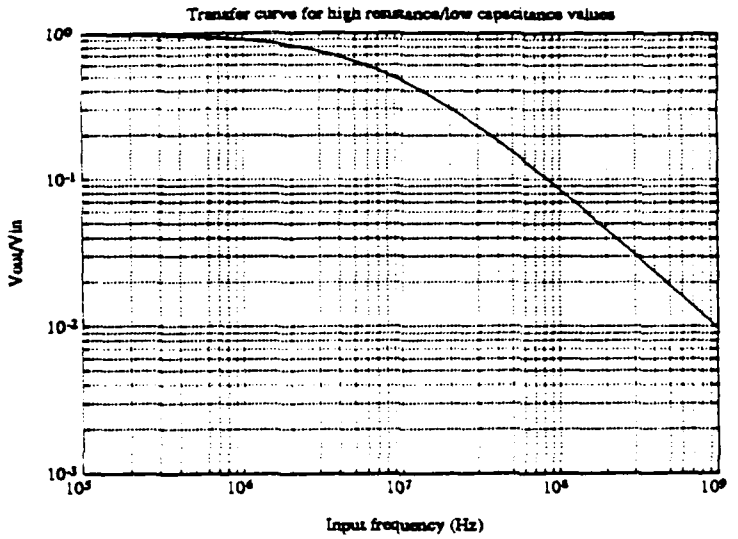

Figure 2.8: Transfer curve for high resistance/low capacitance values.

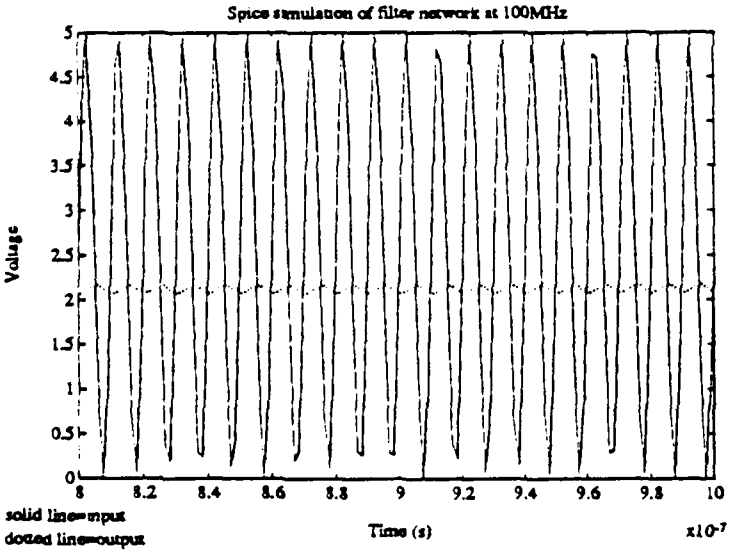

Figure 2.9: Spice simulation of filter driven by 100MHz sinusoid.

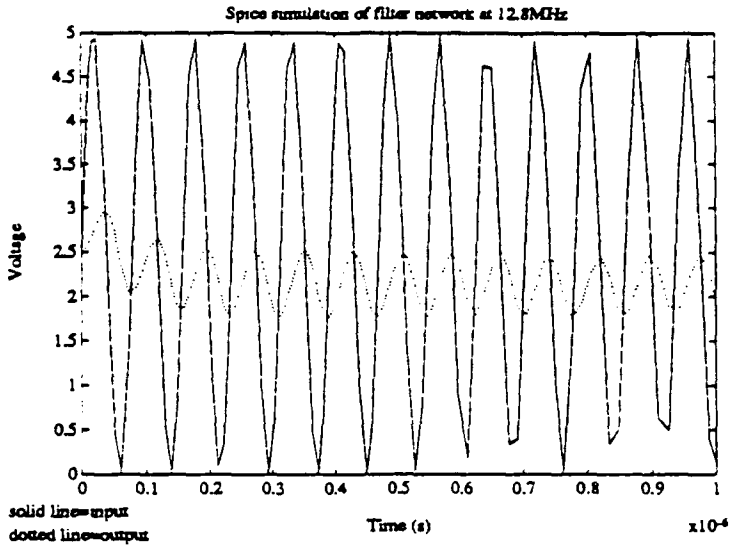

Figure 2.10: Spice simulation of filter driven by 12.8MHz sinusoid.

the "lock-in time" benchmark loses significance. Furthermore, in most applications, it is irrelevent how long the clock takes to lock to the optical signal as long as it does so within a reasonable time.[18] Hence, the low pass/loop filter combination presented here should be adequate for most applications.

2.1.3 The Photodetector and Phase Comparator.

The photodetector/phase comparator/amplifier shown in figure 2.11 represents the most significant modification of the ORCA design presented by Clymer [2] and shown in figure 2.12.

There are several reasons for modifying Clymer's circuit. First, although biasing the photodiode using the output of the VCO does simplify the problem of phase detection, it has implementation problems. The p-n junction used for photodetection in the circuit is comprised of a 20μm by 20μm p^+-diffusion in an

[18] perhaps a fraction of a second.

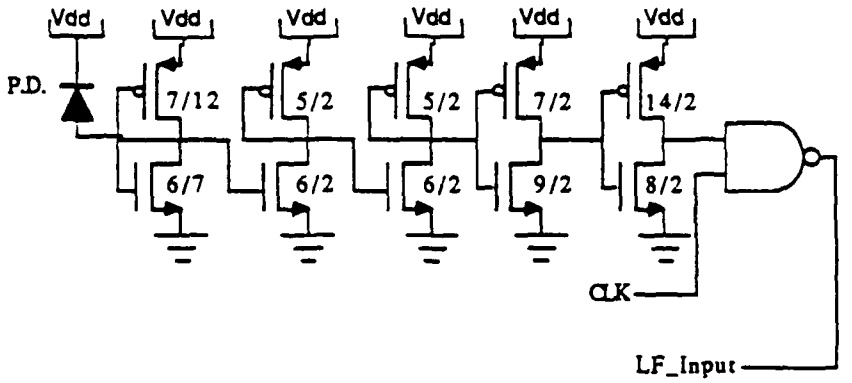

Figure 2.11: Photodetection/amplification/comparator circuit as implemented on OCD2 chip.

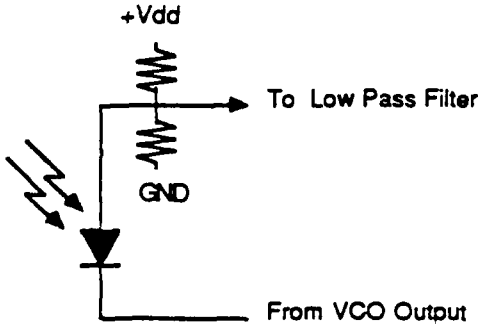

Figure 2.12: Clymer's photodetection and phase comparator circuit

n-well.[19] Due to the size of this diffusion, it has significant capacitance, which causes noticeable loading on the VCO, thereby lowering its operating frequency. To correct this problem, increasing the VCO inverter's sizes is a possibility, but doing so markedly increases the power drawn by the VCO. Also, there is still a problem with this design due to the reverse recovery transient exhibited by the photodiode. In this phenomenon [8, pp. 170–173], the stored junction charge in the p-n diode causes a large reverse current to flow when the VCO output switches from low to high. As an example, for equal pullup and pulldown resistor values of 63kΩ, the reverse transient will have a peak amplitude of 68μA, decaying as the stored junction charge is extracted. Since 68μA is almost seven times as great as the photocurrent we intend to detect, the reverse recovery transient poses a serious problem. Because of these problems we abandon the use of the VCO output to bias the photodetector, instead biasing it with a ratioed inverter with its input tied to its output. To replace the lost phase detection ability of the VCO biased photodiode, we provide a simple digital phase detector later in the circuit.

Another difference between the current design and Clymer's design is the use of an amplifier/inverter chain as opposed to a simple resistor biasing configuration. This difference arises due to technology's constraints, since the resistors in Clymer's design must be on the order of 1MΩ to get the desired range of control voltages.[20] Note one cannot normally get these resistances using biased transistors.[21] Therefore, we resort to detecting the small voltage difference[22] produced by the photodiode, and then amplifying this voltage swing through

[19] To help limit the amount of injected current that reaches surrounding circuitry, a "p-plug" ring is placed around the n-well, "tying" the substrate strongly to ground.

[20] The voltage range of interest is ~3.5–5.0 volts.

[21] Special processes used for memory chips provide a highly resistive polysilicon layer, but the resistances typical of such a layer are on the order of 1GΩ and have large variations in resistance values. Hence, even these special processes are inadequate for this application.

[22] The voltage swing produced by the photodiode is approximately 300mV, according to SPICE simulations. This voltage swing is due to the photocurrent acting through the biasing transistors and is limited by the resistances of the biasing transistors, the parasitic capacitance it must charge, and the operating frequency of the clock.

Photodiode	VCO output	error voltage
Dark	Low	High
Dark	High	High
Illuminated	Low	High
Illuminated	High	Low

Table 2.1: Error voltage as a function of illumination and VCO output.

two analog stages. These identical linear amplifiers are ratioed in such a way that their quiescent point[23] lies in the center of their input voltage range at the anticipated clock frequency, while they maintain a nearly linear voltage transfer characteristic over the input voltage range. In this way, the amplifiers in the chain maintain an average voltage equal to their quiescent voltage, and they are prevented from entering the non-linear region of their voltage transfer characteristics. Once we have amplified the signal sufficiently, we use two high-gain inverters to extract the "logic threshold crossing"[24] information from the signal and to amplify the signal to CMOS logic levels. Finally, the resultant signal is digitally NAND-ed with the VCO output to create an error voltage. The result of the NAND operation is shown in table 2.1.

Spice simulations have been performed to determine what the transistor dimensions for the amplifier chain should be in order to maintain a strong reverse bias on the photodiode, to provide a reasonable gain per amplifier stage, and to keep the voltages involved from leaving the linear region of the amplifier's transfer curve. The resulting transistor dimensions are noted[25] on figure 2.11.

Since this part of the ORCA is mainly analog circuitry, one becomes concerned with the effects of process variations on its functionality. SPICE simula-

[23] Here quiescent point is defined as that voltage to which the amplifier would settle if its output were connected to its input.

[24] i.e. the time when the amplified signal crosses the logic threshold of the first inverter. The logic threshold of the first inverter is set to 2.5V.

[25] Notation is in the form W/L, where W is the drawn gate width in μm and L is the drawn gate length in μm.

tions were done to determine how worst-case[26] uniform variations in transistor geometries[27] and threshold variations[28] would affect the output of the phase comparator at varying values of skew between optical signal and VCO output. Results indicate that at 100MHz, the worst-case[29] "average" output value from the NAND gate occurs for a short and wide channel at 2ns of skew, and this value is approximately 3.4 volts. In order to accomodate this worst case, the VCO is adjusted to run at 100MHz when the control voltage CONTR is approximately 4.2 volts—midway between the 3.4 volts of worst-case NAND output and the 5 volts of "perfect lock" NAND output.

In addition to being concerned over the effect of process variations on the functionality of this circuit, one is also concerned about the limits that this photodetector/amplifier circuit may impose on the phase-locked loop's operating frequency range. SPICE simulations also indicate that if the optical clock frequency is too high, the parasitic capacitance at the input to the first analog amplifier will not charge to a high enough voltage to cause the inverters in the chain to switch fully. Hence, the output signal of the NAND gate will no longer function as a valid error signal. On the other hand, if the optical clock frequency is too low, the parasitic capacitance will charge too much. This causes the analog amplifier stages to saturate and effectively lengthens the pulse seen by the NAND gate, since the inverters detect threshold crossings. Beyond a point, this pulse-widening causes too much loss of skew information between the optical signal and the VCO output, resulting in uncertainty of lock in the loop. (See figure 2.13.) Hence, the ORCA no longer functions properly.[30] Since

[26] Worst case variations are determined by using data acquired from a typical MOSIS 2.0μm CMOS run. See Appendix C.

[27] A reasonable model of the effects of misaligned fabrication masks.

[28] Modelling doping/process variation effects, to first order.

[29] The worst-case is the highest low voltage produced by the unloaded NAND gate, as this is the worst-case bottom of the control voltage range. Five volts is assumed to be the top of the voltage range.

[30] With continued decrease in frequency, the voltages "walk off" the linear portion of the transfer curve. At this point, the input to the NAND gate shows only small (<1 volt) variations and these variations are not near the logic threshold of the NAND. Hence, the NAND loses

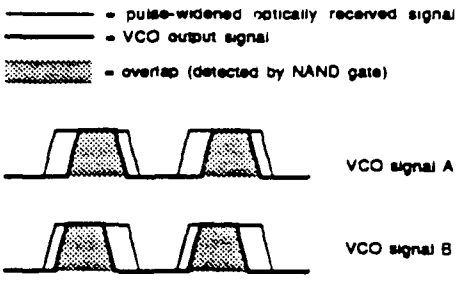

It is evident that the two VCO signals have a difference (skew) with respect to each other. The NAND output is high, except in the overlap region. Since the average NAND output value is the same in both cases, the error voltage on the VCO is the same and therefore the operating frequencies of VCO A and VCO B are the same. However, the pulse widening caused by too low an optical clock frequency allows uncorrectable skew to exist between the two VCO clocks.

Figure 2.13: Illustration of how too low an optical clock frequency can cause loss of skew information.

too short or too long a pulse width causes unreliable operation of the loop, we see that the this photodetector/amplifier/phase-comparator circuit imposes a limit on the frequency range over which the ORCA can lock to an optical input signal. Thus, the loss in frequency range imposed by removing four of the pass transistors in the VCO, as mentioned on page 9, is not a serious limitation.[31]

One final difference between the current design and Clymer's design should be mentioned. In Clymer's design, the output of the phase detector was *high* iff[32] both the VCO output was high and the photodiode was illuminated. In the current design, the output of the phase detector is *low* iff these same VCO and photodiode conditions exist. Hence, whereas Clymer's circuit tended to "pull" the VCO into lock at $0° + \delta$ (where δ is a small steady-state phase error), the current design "pushes" the VCO into lock at $180° + \delta$. This should pose no

all information from the photodetector.

[31] SPICE simulations reveal that the dominant limitation of frequency range of operation arises from the photodetector/amplifier circuit, rather than the VCO.

[32] "Iff" is used in the mathematical convention meaning "if and only if."

problem, since the skew between the optical clock and the VCO is unimportant; it is the skew between VCO's of different ORCAs on the same chip that concerns us. Furthermore, the average output of the NAND ranges from ~3.0–5.0 volts, which is ideal for controlling the VCO. Hence, the NAND output, once filtered, should serve well as the control voltage for the ring oscillator.

Now that the photodetector/amplifier/phase comparator has been described, we examine the circuitry necessary to convert the output signal of the VCO into a useable system clock signal.

Chapter 3

Testing of Chip

One possible optical setup for testing the OCD2 chip, along with appropriate test procedures, is now detailed.

3.1 Optical Test Setup

To facilitate easier testing, only one optical setup should be used for all of the suggested tests. Note that this would help to maintain consistency between tests, since different setups might cause variations in optical power incident on the photodiode, and such variations could affect the operation of the ORCAs.[1] The suggested optical setup in shown in figure 3.1.

As Clymer noted [2, pp. 15–17], a visible (i.e. He-Ne) laser source would have higher efficiency since shorter wavelength light results in capture of more photons in the photodiode's depletion layer. However, most visible lasers must be externally modulated by an acousto-optic light modulator (AOLM). Since our circuit expects 5ns pulses (100MHz square wave), we cannot afford a rise or fall time of more than about 1ns. Thus, we would require an AOLM with about

[1] This is due to the previously mentioned sensitivity of the photodiode's amplifier chain. See page 20.

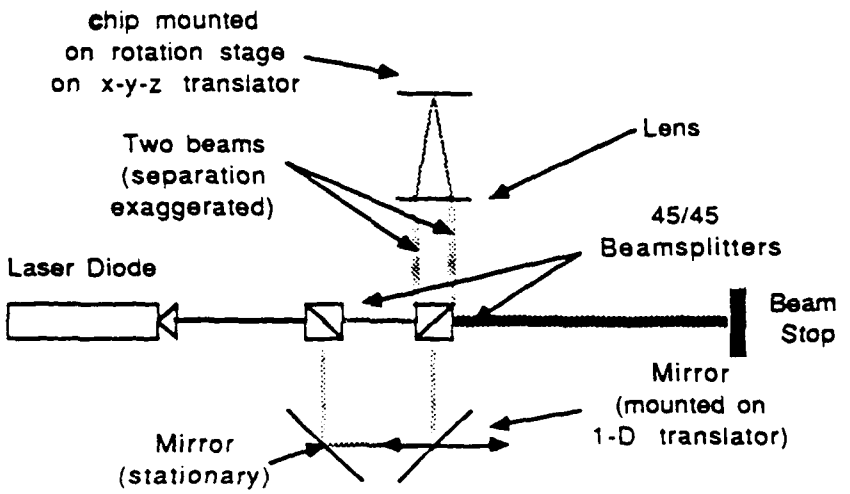

Figure 3.1: Optical setup used in all chip tests.

1GHz of bandwidth. Such AOLMs are extremely expensive[2] and are not very efficient. As an alternative, a laser diode system can be directly modulated, but most laser diodes operate in the infrared (at approximately 820nm) where the p-n photodiode is not very efficient. However, one laser diode system was found which features a 750nm beam, a modulation input, and colliminating optics. Thus, the capture efficiency is better than for most laser diode systems, and we needn't worry about aligning the diode to collimating optics. Furthermore, its specifications (which are included in Appendix D) indicate that the modulation input has a range of 100kHz–500MHz. Since 500MHz corresponds to 1ns each of rise and fall times, this laser diode system could serve our purposes very well.

In one possible setup, the laser beam emitted from this system would follow the path indicated in figure 3.1. Note that each beam is reflected once and transmitted once, so each should have the same optical power (neglecting *mirror losses*) *even if the beamsplitters* are slightly anomalous in their trans-

[2] on the order of $20,000.

mittance/reflectance characteristics at this wavelength. (Sample beamsplitter specification sheets are also provided in Appendix D.)

After leaving the beamsplitter cubes, the beams would then be focused by the lens into two small spots on the chip. The chip can be translated in all three dimensions and can be rotated in order to position the two beams so that they would impinge on two adjacent ORCAs' photodetectors while allowing control of the spot size. The separation of the two beams could be controlled using the 1-D translator on the second mirror.

Optical power considerations for this suggested setup are as follows: the suggested laser emits up to 4mW of optical power. Each sample beamsplitter is 45% transmitting/45% reflecting. Hence, each beam of light has 810μW of power exiting the second beamsplitter. (We neglect losses due to the mirrors.) With a suggested lens[3] we could focus the 4.5mm x 2.0mm spot down to a 27μm x 12μm spot. However, to make sure the photodetector is fully illuminated, we would defocus the beams so that the smallest dimension is 30μm. At this point, the optical power incident on each detector is approximately 324μW.[4] Since the responsivity of the p-n junction is approximately 0.12–0.15 A/W (neglecting reflection losses at the chip), [5, 9, pp. 743–756], this corresponds to a photocurrent of at least 39.2μA. We adjust the power and/or focusing of the beam to scale this photocurrent down to 10μA, since this is the value we used in the SPICE simulations.

We now proceed to detail suggested tests, utilizing this suggested optical setup.

[3]The lens has a focal length of 70mm and a diameter of 60mm.

[4]Since the photodetector is square and the incident beam in elliptical in nature, the rest of the optical power would be lost on the shielding around the photodetector.

3.2 Pre-testing

Before attempting to test any given chip using the previously described optical test, one should verify the functionality of as much of the chip as possible without the optics.

To test the skewdetector blocks, the clocks should be allowed to run at whatever natural frequency they settle to in the absence of an optical signal. An oscilloscope should be used to monitor the charging characteristics of the skew pins (Skew0–Skew11) while periodically activating the discharge line. Also, the same pins should be monitored while the calibrate line is active. In this way, "stuck at zero" and "stuck at one" faults on the skew pads and the discharge pad could be detected.

For the frequency divider testing, the four frequency output pins (FreqOut0–FreqOut3) should be monitored while the various select lines (Select3–Select0) are set in turn to each hexadecimal value from 0 to 8. An oscillation on the output pins with a frequency in the "reasonable" range indicates that the corresponding phase of the corresponding clock is oscillating correctly and that the multiplexor/frequency divider circuit is operational at the natural clock frequency. Note that the "Init" line needs to be pulsed before the frequency divider is guaranteed to function properly.

For the shift register testing, one should monitor the four error pins (Error3–Error0). The Init line should be pulsed to set the T flip-flop to a guaranteed operational state, and after a short while (roughly 100 clock cycles) the shift register should have loaded with alternating ones and zeros. From this point on, there should be no toggling of the error lines. Observation of the turn-on transient of this line would verify the lack of "stuck at" faults in the shift register circuit, and observation of this line after the shifter is loaded would verify that the shift register test is operating correctly.

This completes the testing which can be done without the use of the optical setup. We now proceed to the suggested tests which can be performed using

the optics.

3.3 Testing with Optics

Those chips which pass the suggested pre-testing should then be inserted in the optical test setup. After adjusting the beams so that approximately 10μA of photocurrent is generated in each of two adjacent photodetectors, the following measurements should be made.

3.3.1 Skew Testing

With the calibrate line inactive, a square wave should be applied to the initialization line to provide for periodic discharging of the skew lines. By studying the charging characteristics of the output line responsible for measurement of skew between the two ORCAs of interest and comparing these characteristics to those generated by the same output line when the calibrate line is active, a skew measurement could be obtained. This same measurement should then be made with the introduced path length difference affecting the second ORCA's photodetector. An average skew should then be computed.

3.3.2 Frequency Output

Unfortunately, to drive the output pins *reliably*, one must frequency divide the oscillator before outputting it. This results in a loss of frequency resolution. However, in this application, we merely wish to verify that the oscillator has locked to the optical clock frequency and not some harmonic of it. To do this, we set the select lines (Select3–Select0) to the hex value of the oscillator under consideration and observe the output on the frequency output pins (FreqOut3–FreqOut0).

3.3.3 Shift Register Test

In this test, the "Error" lines are observed and an estimate of how many transitions it makes in a ten second interval[5] is taken. Note that each transistion (excluding those involved in the start-up transient) indicates that an error occurred.

[5]Note: the start-up transient should not included in the ten-second interval of test. In this way, there should ideally be no transistions during testing.

Chapter 4

Conclusions

Here we present a summary of the work done and suggestions for future work.

4.1 Summary of Contribution

This work demonstrates that a reasonably modular circuit can be designed to implement a phase-locked loop approach to optical clock distribution. Simulations indicate that said circuit can be operated at a clock frequency of 100 MHz, and a test chip featuring this circuit along with appropriate diagnostic circuitry has been fabricated. Finally, suggested test setups and procedures have been detailed.

4.2 Future Work

The most obvious suggestion for future work is to perform the indicated tests. Assuming that said tests confirm the functionality of the circuit, a more elaborate circuit should be designed—perhaps implementing a non-trivial high speed function such as a discrete Fourier transform—using this optical clock distribution circuit. Such a device would require a far more sophisticated optical

system including a hologram to focus the optical clock onto the receptors, a higher power laser to drive more than two receptors at a time, and a high speed optical modulator.

On the other hand, should the fabricated chips fail their diagnostic tests, it would be necessary to determine why. Perhaps current SPICE models are simply not adequately accurate for this process. Perhaps skew introduced by detection of the light is excessive and non-standard processes would need to be utilized. There are indeed a number of directions to explore given a failure of the test chips to perform.

Appendix A

Pad Assignments

A summary of the pad assignment on the OCD2 chip.

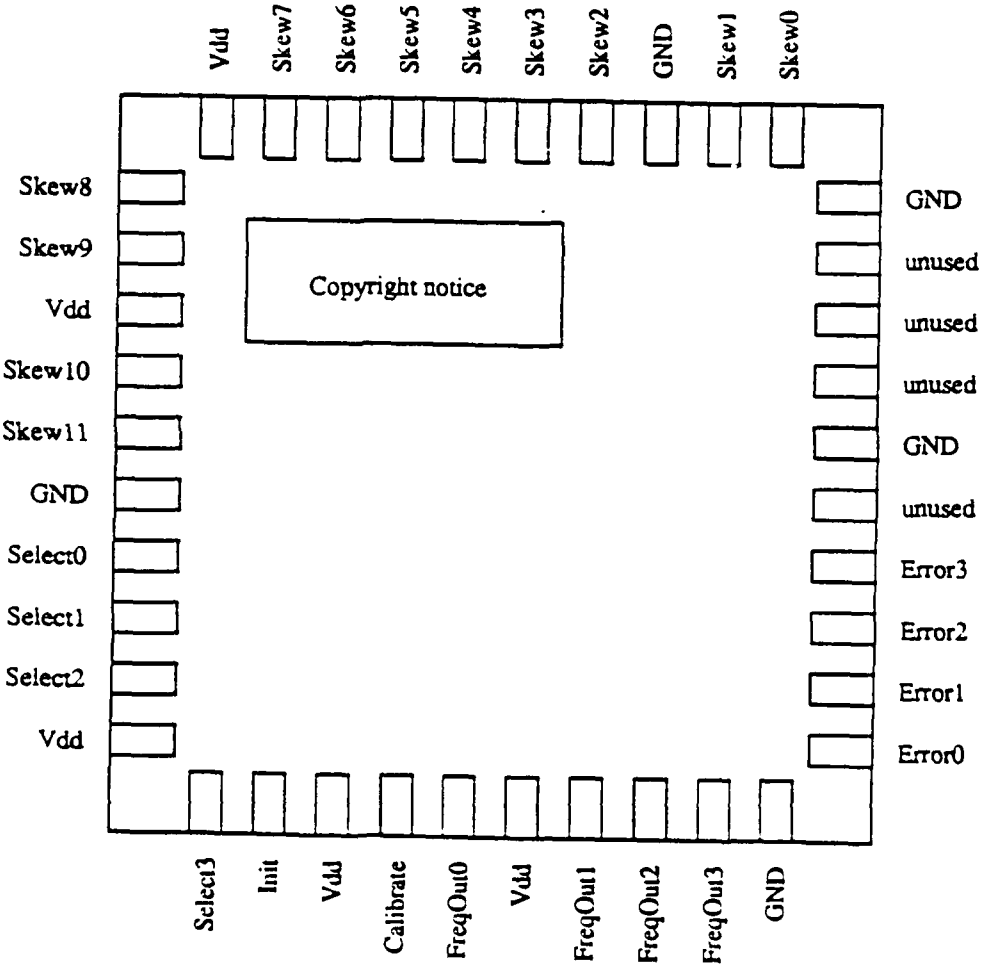
Vdd
Skew7
Skew6
Skew5
Skew4
Skew3
Skew2
GND
Skew1
Skew0
Skew8
Skew9
Vdd
Skew10
Skew11
GND
Select0
Select1
Select2
Vdd
Copyright notice
GND
unused
unused
unused
GND
unused
Error3
Error2
Error1
Error0
Select3
Init
Vdd
Calibrate
FreqOut0
Vdd
FreqOut1
FreqOut2
FreqOut3
GND

Appendix B

SPICE parameters

A summary of SPICE3 parameters used in simulating the analog portions of the OCD2 chip.

SPICE parameter	NMOS value	PMOS value
level	2	2
vto (V)	0.812464	0.747301
tox (m)	413E-10	413E-10
nsub (cm^{-3})	8.644E+15	1.774E+16
xj (μ)	0.4	0.4
ld (μ)	0.065	0.071
u_0 (cm^2/V-s)	535.36	240.17
ucrit (V/cm)	39202.4	10000
uexp	7.6987E-2	0.10088
vmax (m/s)	5.1E+4	3.5067E+4
delta	0.36	1.0E-6
rsh (Ω/sq)	34.37	129.5
cgso (F/m)	1E-10	1.5E-10
cgdo (F/m)	1E-10	1.5E-10
cj (F/m^2)	1.518E-4	2.535E-4
cjsw (F/m)	5.87E-10	3.33E-10
mj	0.6744	0.5213
mjsw	0.3087	0.2801
pb (V)	0.71	0.71
kp (A/V^2)	4.476e-5	2.008E-5
lambda (V^{-1})	3.13E-2	2.46E-2
gamma ($V^{1/2}$)	0.3698	0.5298
phi (V)	0.6	0.6
nfs (cm^{-2})	5.05E12	3.19697E12
js (A/m^2)	1E-14	1E-14
neff	1.001E-2	1.001E-2
nss (cm^{-2})	0	0
tpg	1	-1

Appendix C

Process variations

Selected measured process variations for the 2.0μm n-well CMOS MOSIS run M71K are presented. These are the worst case variations and were used in the simulations of the amplifier chain on page 19.

process parameter	worst case low	worst case high
NMOS threshold voltage (V)	0.6424	1.265
PMOS threshold voltage (V)	-0.7038	-1.354
NMOS delta length (μm)	-.352	+.352
PMOS delta length (μm)	-.216	+.216
NMOS delta width (μm)	-.577	+.577
PMOS delta width (μm)	-.732	+.732

Note that the delta length and delta width are the variations with respect to some nominal length or width. For example, a 2μm NMOS drawn channel length may actually vary from 1.648μm to 2.352μm.

Appendix D

Specifications

Copies of the specification sheets for the following test equipment used are provided on the following pages.

1. Laser Diode specifications taken from *Oriel Lasers & Accessories Catalog*, © 1988 by Oriel Corporation, Stratford, CT.

2. Beamsplitter specifications taken from *Melles Griot Optics Guide 4*, © 1988 by Melles Griot. Used with permission.

ORIEL DIODE LASER SYSTEMS

- Compact, Economical, Complete Packages
- Choice of Wavelength and Power
- Modulation Input
- Powers to 25 mW
- Various Beam Sizes
- Meet CDRH Safety Requirements

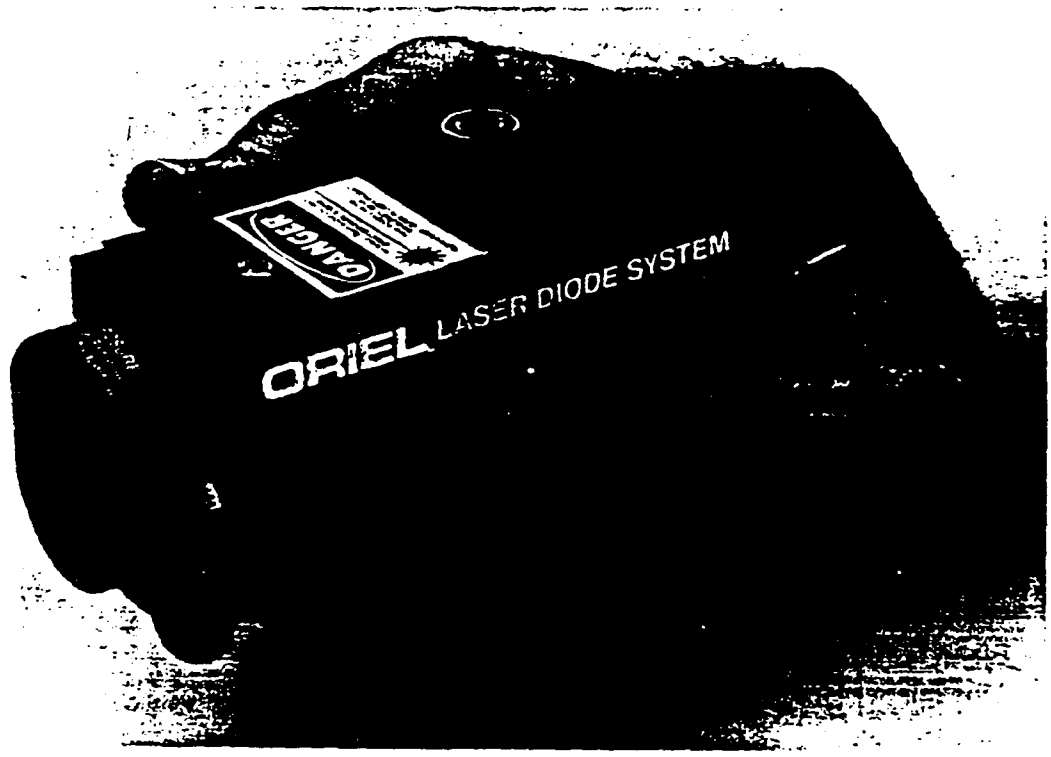

Oriel's new compact Diode Laser Systems are ideal for the development of applications of these versatile solid state lasers. Our units are complete, ready to plug in and turn on. The tiny laser diodes, which are vulnerable to static, are already safely wired. Pre-aligned optics collect and collimate the light efficiently.

ORIEL DIODE LASER SYSTEMS

Oriel Diode Laser Systems use proven GaAlAs lasers with output powers from 4 to 25 mW and wavelengths from 750 to 830 nm. The lasers are built on tiny chips of p-type GaAs substrates. (Fig. 1)

The active layer is only a fraction of a micron thick. When current flows through this layer, light is emitted as electrons and holes recombine. The refractive index of the active layer is higher than that of the layers above and below.

Any light striking the interface at a large angle of incidence is internally reflected. As in a fiber optic the light is guided in the high refractive index region. A V groove through a current blocking layer confines the current flow laterally, defining the gain region. Light which escapes is absorbed in the unpumped crystal. This is called gain guiding. The dimensions of the lasing region ensure stable operation of the fundamental transverse mode. No other modes can operate. The higher power chips have optical coatings on the front and rear surfaces to optimize light extraction.

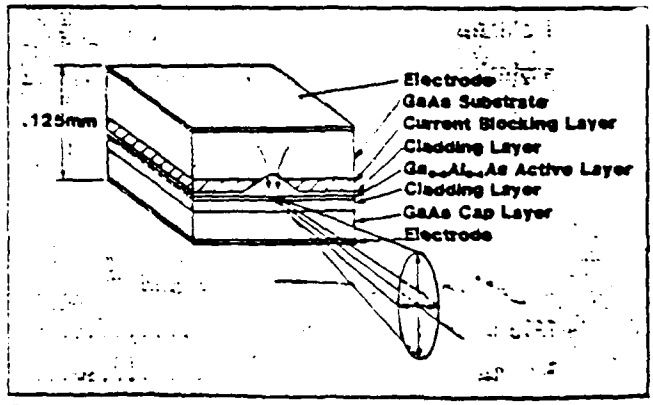

Fig. 1

Each laser chip is mounted in a can as shown in Fig. 2. A PIN photodiode located behind the laser chip continuously monitors the output from the rear of the laser. We use the photodiode signal in a feedback loop to stabilize the laser output. A large aluminum heatsink keeps the diode can cool. Our drive circuit powers the laser and maintains steady output. The circuit includes a DC loop for control of total power and an AC loop to minimize noise. Bypass capacitors and a slow start circuit protect the laser chip from fatal surges.

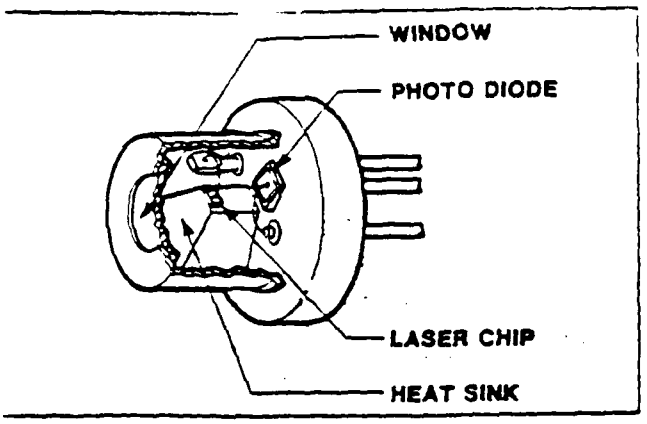

Fig. 2 The laser and monitor photodiode.

LASER CONTROLS

Operating any Oriel Diode Laser System is straightforward. Simply plug the DC power supply into the rear of the laser head and into the line, wire the safety interlock appropriately, and turn the key switch. A green warning light is activated and a few seconds later the laser comes on. A control knob on top of the laser changes the output power from about 20 to more than 100 % of rated power. The safety shutter may be used to stop the beam.

A BNC jack at the rear of the laser accepts modulation input.

MODULATION INPUT

You can modulate Oriel Diode Laser Systems at frequencies from 100 kHz to 500 MHz by using the BNC connector on the rear of the laser. The input is AC coupled. Lower frequency modulation is not possible except to special order. Typically 3-5 volts into the 50 ohms input results in 50 -60 % modulation. Up to 90% modulation is possible with 17 dBm RF power. The power control circuit automatically adjusts DC bias power to keep the average output power constant.

LIFETIME

Diode lasers usually fail for one of two reasons.

- Thermal stress on the semiconductor.
- Damage to the laser output facet.

Thermal stress from running the diode (or from soldering it in place) can result in the growth of tiny flaws in the crystal structure. When crystal defects become large, the diode cannot lase.

As defect development and propagation is temperature dependant, the lifetime is dramatically shortened by running at higher than design temperatures. MTBF drops from 200,000 hours for a case temperature of 30 degrees to 20,000 hours (2.5 years) for a case temperature of 65 degrees. The higher power lasers have shorter lifetimes.

The laser power density at the output facet of the chip is very high. This leads to gradual darkening of the facet from photochemical changes. The MTBF figures indicate that this is not a serious problem. However, static, power surges, or high turn on currents can result in momentary high power operation of the laser. At 25 mW on a 1 by 5 micron typical output facet, the continuous power density in normal operation exceeds several MW/cm². Much higher power density from transients causes catastrophic damage to the chip surfaces which act as the laser reflectors, destroying the laser.

SAFETY

Appropriate safety measures should be used when working with any laser. The high output power from the tiny diode laser package, and the fact that most are invisible, demands special precautions for the operator and control of access to the working area. Some of our diode laser beams are barely visible. This can be misleading. A 25 mW laser producing a faint red spot is as dangerous as a bright 25 mW argon ion laser, and potentially more so because of the natural association of danger with the bright beam and the protection from the "blink response".

Our lasers include a shutter, key switch, labels and interlock as required. Safety eyewear is listed on pages 30 to 31 and door signs on page 32.

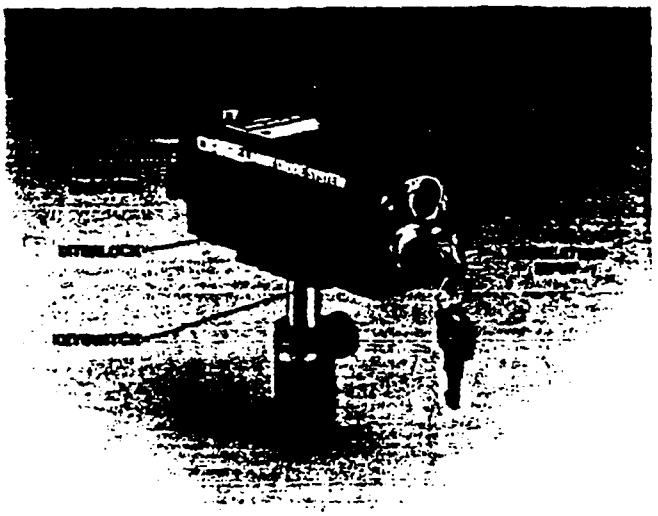

Fig. 3

BEAM CHARACTERISTICS

Diode lasers do not have the familiar circular beam shape with Gaussian intensity distribution of a TEM00 Helium Neon Laser. The beam is elliptical, does not have a true Gaussian profile, but does have the high brightness and low divergence necessary for focussing to small areas or coupling to optical fibers.

Our lasers are all fundamental transverse mode TE00. The assymetry in the vertical and horizontal lasing region appears in the collimated elliptical beam which has low but different divergences parallel and transverse to the lasing stripe.

The beam is also astigmatic like most diode laser sources. The minor and major axes of the ellipses appear to have originated in different spots a few microns apart. Refocussed beams have the minima of the major and minor axes in different positions. This is unimportant except for critical tight focussing applications.

Fig. 6 shows actual beam shapes of our 79426 Laser System. This is typical for our devices with larger beams. The beam characteristics of the smaller beam systems are more complex. The source wavefronts and abberations of the very high F/# lens used to collimate the laser light from this source result in unusual beam patterns up to 500 mm from the source. (See page 49 for more beam profiles).

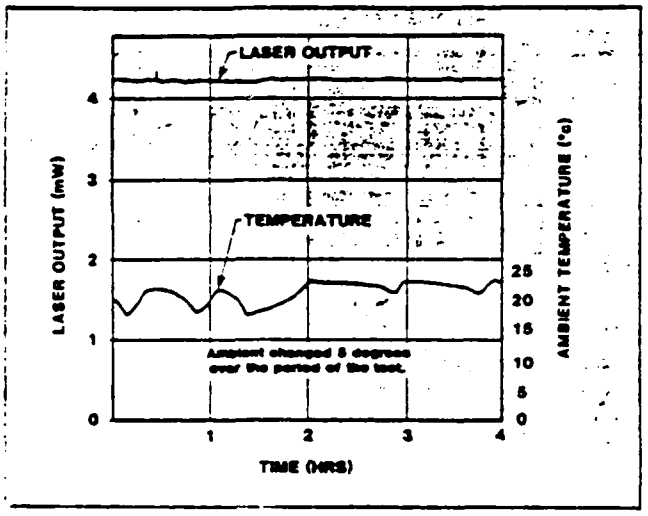

Fig. 4 Output of the Model 79426 over a 4 hour period.

Fig. 5 Typical short term power variation of the model 79426.

OUTPUT STABILITY

The laser output is constantly monitored by the photodiode which is not sensitive to temperature. The output is temperature compensated. The slow DC loop stabilizes the power to better than 2% (Fig. 4). The A.C. loop maximizes signal to noise ratio. (Fig 5).

The 75 mW lasers can exhibit occasional power fluctuations of a few percent due to longitudinal mode transitions. All these lasers are sensitive to reflection of the beam back into the diode. This can cause completely unstable operation.

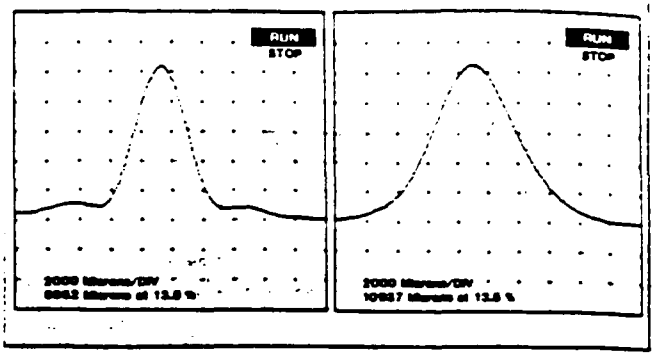

Fig. 6 Horizontal (left) and vertical (right) intensity profiles of the 79426 Laser System at 4m from the laser output.

SPECTRUM

We offer diode laser systems with a choice of wavelength from 750 to 830 nm. The actual wavelength will be within +/- 5 nm of the nominal. When the laser is warmed up and running at full power, the output spectrum is the sharp spike characteristic of single longitudinal mode operation. The width of this spike was measured as 0.6 GHz (1.2 pm) for our 780 nm 5 mW laser. The linewidth was measured using a 7.5 Ghz FSR Scanning Interferometer. (Fig. 7) The narrow line has a sensitivity to diode can temperature of about 0.25nm/ degree C. Significant variations in ambient will influence the wavelength but not in a smooth monotonic fashion. Mode hops of a fraction of a nanometer occur as the wavelength increases with temperature.

The wavelength also changes with the power control. (Fig. 8, page 49) At low power, some lasers operate in many longitudinal modes and the spectrum appears as the familiar comb of modes. To ensure single mode output, operate these devices at maximum power.

POLARIZATION

The laser chips are polarized parallel to the minor axis of the ellipse. The polarization of the beams from our systems is nominally horizontal with the laser system supported normally. When the power is measured, while rotating one of our crystal polarizers in the beam, the ratio of maximum to minimum is typically 75:1.

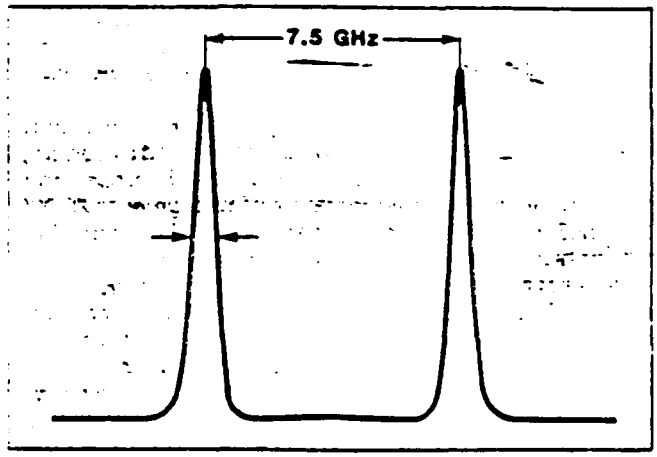

Fig. 7 Spectrum of the Model 79403 laser recorded with a scanning Fabry Perot Interferometer. The free spectral range of the interferometer was 7.5 GHz and the laser was running at full power.

OUNTING

The laser can be easily rod mounted for use on optical ples or benches. There are three tapped 1/4-20 holes in the ser base. These holes are in a line parrallel and under the tical axis. The holes are spaced 1.5 inches (38 mm) apart. r extensive range of rods, rod holders and accessories n be used to hold the laser in any orientation. See Volume I these products.

OUPLING THE OUTPUT TO OPTICAL FIBERS

The output of the small beam lasers may be directly upled to an optical fiber using the coupling systems scribed on page 76. These fit directly on the threaded tput of our laser systems. These compact couplers can o be mounted in the 77842 holder for use after the beam s exited the laser system. In either case be sure that the upler chosen matches the beam and fiber sizes. For ernative coupling arrangements see page 76.

RODUCT LISTING AND RDERING INFORMATION.

odel o.	Wavelength nm	Power mW	Beam Size mm X mm	Beam Div* mrad	Price ($)
3401	750	4	1.5 X 0.7	<1	$785.00
3402	750	4	4.5 X 2.0	<0.5	$785.00
3403	750	4	8.9 X 4.0	<0.4	$785.00
3406	780	4	1.5 X 0.7	<1	$675.00
3407	780	4	4.5 X 2.0	<0.5	$675.00
3408	780	4	8.9 X 4.0	<0.4	$675.00
3411	780	8	1.5 X 0.7	<1	$995.00
3412	780	8	4.5 X 2.0	<0.5	$995.00
3413	780	8	8.9 X 4.0	<0.4	$995.00
3416	780	16	1.5 X 0.7	<1	$1325.00
3417	780	16	4.5 X 2.0	<0.5	$1325.00
3418	780	16	8.9 X 4.0	<0.4	$1325.00
3421	810	4	1.5 X 0.7	<1	$785.00
3422	810	4	4.5 X 2.0	<0.5	$785.00
3423	810	4	8.9 X 4.0	<0.4	$785.00
3428	830	25	1.5 X 0.7	<1	$1625.00
3427	830	25	4.5 X 2.0	<0.5	$1625.00
3428	830	25	8.9 X 4.0	<0.4	$1625.00

e 50% intensity points are measured at 1 m and 5 m from the stem output.

ECIFICATIONS COMMON TO ALL ORIEL DE LASER SYSTEMS

ut mode:	Single Transverse Single Longitudinal at full power
spacing:	Approximately 0.32nm
zation Ratio:	Better than 50:1
ng Stability:	Better than 0.1 mrad
ed Ambient:	-10 to 50°C

Dimensions of the Oriel Diode Laser Systems

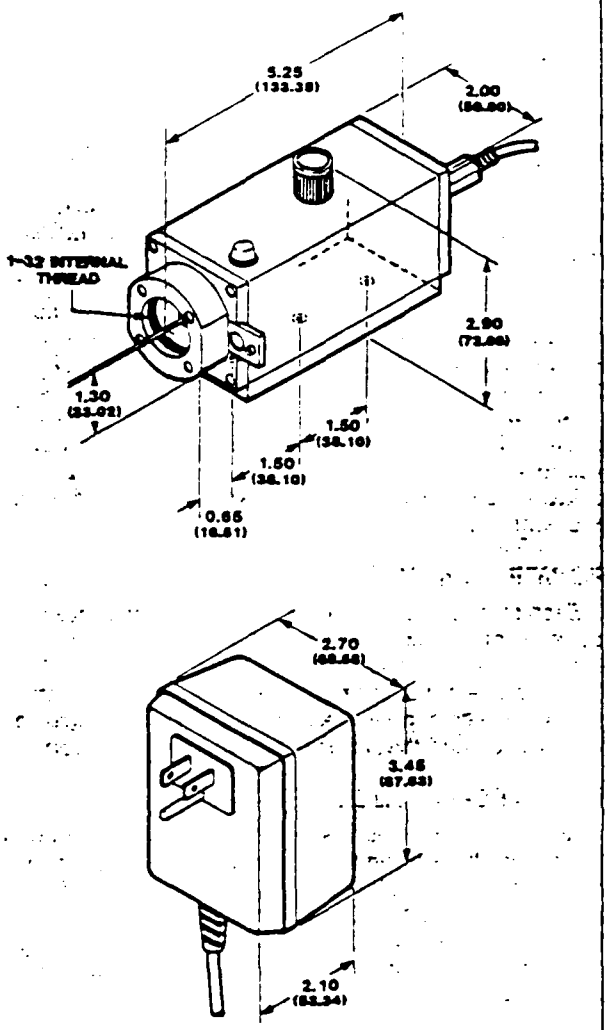

Dimensions in inches (mm).

CDRH warning logotypes, similar to that shown above, appear on each laser to indicate the CDRH classification and to certify that the output power of the laser will not exceed the power level printed on the logotype.

...out 10%, and reflection and absorption are both approximately ...%, with the s- and p-polarized components within 10% of ...each other. Over a wide range of wavelengths, the s- and p-...olarized components stay within 10% of each other. The broad-...and spectral flatness of these beamsplitters makes them ideal for ...avelength scanning instruments.

Another advantage of these hybrid coatings is that they are ...rly insensitive to changes in the angle of incidence. Although ...e cube beamsplitter configuration usually causes the angle of ...idence to be 45°, this angular invariance in the performance ...s the advantage of making these beamsplitters chromatically ...utral for convergent or divergent beams.

Melles Griot hybrid coatings cover the visible spectrum (400-...nm) or the very near-infrared for laser diode use (700-...nm). Longer or shorter wavelength versions of this coating ...e not possible due to severe material limitations.

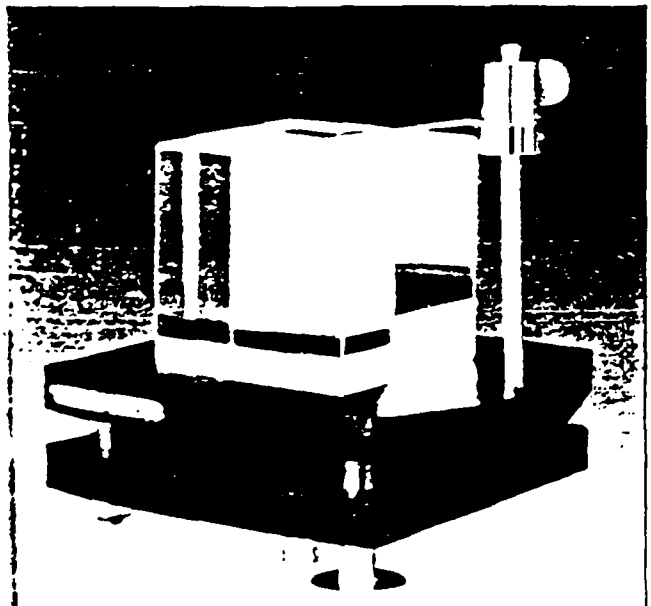

BROADBAND PARTIAL REFLECTION COATINGS

Broadband partial reflection coatings provide a high degree of ...ess. There is negligible absorption in the coating, and the ...ted and transmitted components are almost equal over a broadband wavelength range, averaged over plane s-polarization. However, they are extremely polarization sensitive, with the s- and p-components differing by as much as 70%. Great care should therefore be taken when using these broadband beamsplitters to consider polarization implications for the optical system into which they are to be integrated.

Melles Griot offers broadband partial reflection coatings for the visible spectrum (450-650nm), the laser diode short-waveband (650-900nm), the neodymium laser region (900-1300nm), and the laser diode telecommunications waveband (1300-1600nm).

ALL-DIELECTRIC NON-POLARIZING COATINGS

This range of cube beamsplitter coatings is intended for those applications where polarization effects must be kept to an absolute minimum. Unlike the hybrid coatings, these dielectric coatings are designed for high performance at specific wavelengths, where they easily exceed the performance of any other available beamsplitters. Being totally dielectric, they have negligible absorption.

At the design wavelength, each of these beamsplitters reflect 50 ± 5% of incident light. The s- and p-components of the reflected (and therefore transmitted) beam differ by less than 5%, (i.e. each is within ±2.5% of the average polarization performance).

The current range of these beamsplitters covers two important laser wavelengths, but other wavelengths can be accommodated on special request. The laser wavelengths are the red helium neon line at 633nm, and the 780nm laser diode wavelength.

ANTIREFLECTION COATINGS

All our cube beamsplitters are antireflection coated on all four faces to minimize ghost images and reflection losses. They are normally supplied with a multilayer HEBBAR™ antireflection coating. There is no need to append a Coating Suffix.

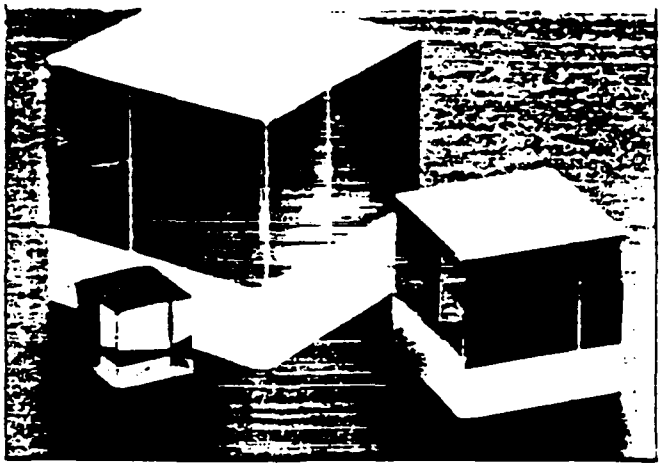

CUBE BEAMSPLITTERS

Cube beamsplitters have several advantages over plate beamsplitters and are widely used for the following reasons. These are rugged beamsplitters which are easy to mount and are ideal for beam superposition applications. This type of beamsplitter deforms much less when subjected to mechanical stress than does a plate beamsplitter. Most of the unwanted reflections from a cube beamsplitter are in the retro direction and do not contribute to ghost images. The coating is very resistant to degradation with time since it is sealed within the body of the cube.

If cube beamsplitters are used in convergent or divergent portions of an optical beam, they will contribute substantial amounts of unwanted aberration. This can be avoided or minimized by only using these components with collimated, or nearly collimated, beams. Conjugate distances which include cubes therefore should be long. Alternately, other elements of the system can be designed to compensate for any aberrations introduced by the cube in a non-collimated beam.

Cube beamsplitters consist of matched pairs of identical right-angle prisms, with their hypotenuse faces cemented together. Prior to cementing, a partial reflection film is deposited onto one of the hypotenuse faces. The prism which is coated is marked with a small dot on one of the ground faces. For best results, the incident beam should be on one of the faces of this prism

We currently offer three types of cube beamsplitter, classi by the nature and performance of the partial reflection coa used. The metal-dielectric hybrid coatings reflect and tran equal portions of incident light with 10% absorption and fairly small differences in s- and p-polarization over an exten wavelength range. The broadband, all-dielectric coatings ab negligible amounts of the incident intensity, but the s- an polarized performance characteristics are quite different. In a tion, Melles Griot now offers a range of non-polariz all-dielectric, partial reflection coatings for these cube beamsp ters which have almost identical s- and p-reflectances at spec laser wavelengths.

Each of the coating types has its particular merits and l tations, so your selection should depend on the inten application.

HYBRID PARTIAL REFLECTION COATINGS

Hybrid, metal-dielectric coatings combine the benefits of b metals and dielectrics to produce a moderate absorption be splitter with little polarization sensitivity. Typically, absorpoo

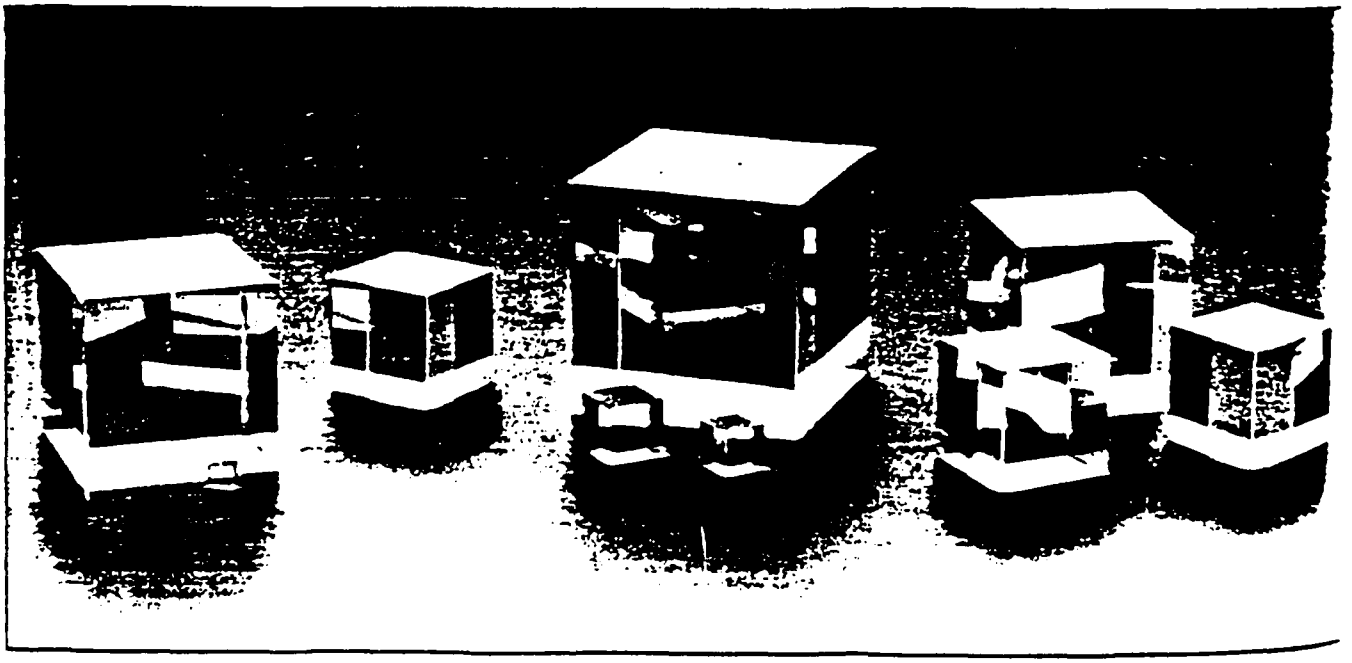

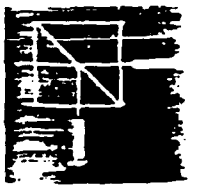

...out 10%, and reflection and absorption are both approximately ...%, with the s- and p-polarized components within 10% of each other. Over a wide range of wavelengths, the s- and p-polarized components stay within 10% of each other. The broadband spectral flatness of these beamsplitters makes them ideal for wavelength scanning instruments.

Another advantage of these hybrid coatings is that they are fairly insensitive to changes in the angle of incidence. Although the cube beamsplitter configuration usually causes the angle of incidence to be 45°, this angular invariance in the performance has the advantage of making these beamsplitters chromatically neutral for convergent or divergent beams.

Melles Griot hybrid coatings cover the visible spectrum (400-...0nm) or the very near-infrared for laser diode use (700-...00nm). Longer or shorter wavelength versions of this coating are not possible due to severe material limitations.

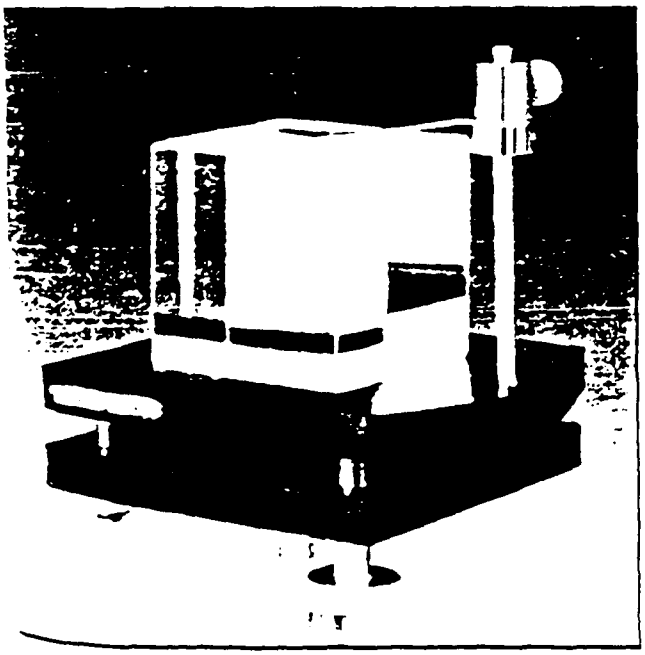

BROADBAND PARTIAL REFLECTION COATINGS

Broadband partial reflection coatings provide a high degree of ...ency. There is negligible absorption in the coating, and the ...ted and transmitted components are almost equal over a broadband wavelength range, averaged over plane s-polarization. However, they are extremely polarization sensitive, with the s- and p-components differing by as much as 70%. Great care should therefore be taken when using these broadband beamsplitters to consider polarization implications for the optical system into which they are to be integrated.

Melles Griot offers broadband partial reflection coatings for the visible spectrum (450-650nm), the laser diode short-waveband (650-900nm), the neodymium laser region (900-1300nm), and the laser diode telecommunications waveband (1300-1600nm).

ALL-DIELECTRIC NON-POLARIZING COATINGS

This range of cube beamsplitter coatings is intended for those applications where polarization effects must be kept to an absolute minimum. Unlike the hybrid coatings, these dielectric coatings are designed for high performance at specific wavelengths, where they easily exceed the performance of any other available beamsplitters. Being totally dielectric, they have negligible absorption.

At the design wavelength, each of these beamsplitters reflect 50 ± 5% of incident light. The s- and p-components of the reflected (and therefore transmitted) beam differ by less than 5%, (i.e. each is within ±2.5% of the average polarization performance).

The current range of these beamsplitters covers two important laser wavelengths, but other wavelengths can be accommodated on special request. The laser wavelengths are the red helium neon line at 633nm, and the 780nm laser diode wavelength.

ANTIREFLECTION COATINGS

All our cube beamsplitters are antireflection coated on all four faces to minimize ghost images and reflection losses. They are normally supplied with a multilayer HEBBAR™ antireflection coating. There is no need to append a Coating Suffix.

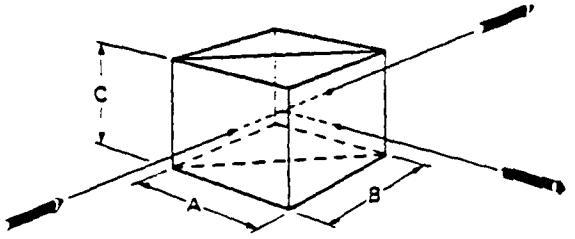

03 BSC

SPECIFICATIONS: CUBE BEAMSPLITTERS

Dimensions: ±0.3mm
Material: BK 7 grade A fine annealed
Face Flatness: 2λ per clear aperture at 546nm
Transmission
- Hybrid: 45 ± 6% (±5% variation with wavelength)
- Broadband Dielectric: 50 ± 5% (±3% variation with wavelength), for average polarization
- Non-Polarizing Dielectric: 50 ± 5% for any polarization with the s- and p-components matched to within 3%

Absorption
- Hybrid: 10%
- Broadband Dielectric: <0.5%
- Non-Polarizing Dielectric: <0.5%

Coatings: All 4 faces HEBBAR™ antireflection coated
Cosmetic Surface Quality: 60-40 scratch and dig

Melles Griot's prism tables and post-mounted prism holders are ideal for holding these cube beamsplitters in an adjustable yet stable manner. See the Components Holders section of this guide for a complete description listing of these useful mounts.

Cube Beamsplitters: Visible Wavelength Range

A=B=C	PRODUCT NUMBER	
(mm)	Hybrid Coating 400-700nm	Broadband 45
5.0	03 BSC 001	03
10.0	03 BSC 003	03
12.7	03 BSC 005	03
20.0	03 BSC 007	03
25.4	03 BSC 009	03
30.0	03 BSC 011	03
40.0	03 BSC 013	03
50.8	03 BSC 015	03

Non-Polarizing Cube Beamsplitters

A=B=C	PRODUCT NUMBER	
(mm)	633nm Coating	780nm
5.0	03 BSL 042	03
10.0	03 BSL 043	03
20.0	03 BSL 044	03
25.4	03 BSL 045	03

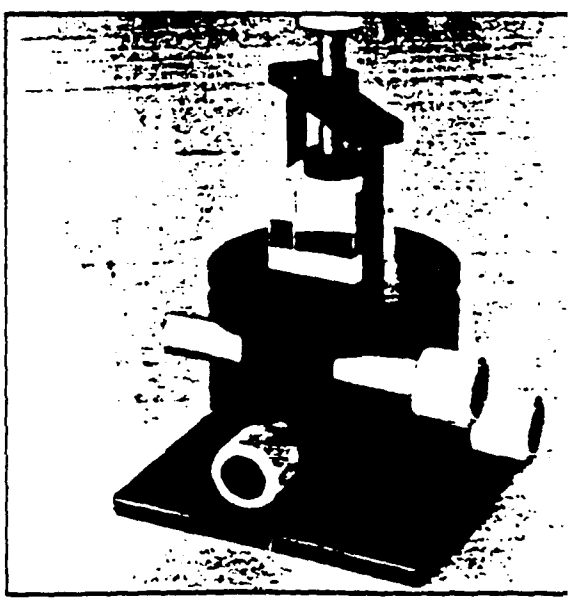

eamsplitters: Near-IR Wavelength Range

PRODUCT NUMBER			
Hybrid Coating	Broadband Coating		
700-1100nm	650-900nm	900-1300nm	300-1600nm
03 BSC 023	03 BSD 042	03 BSD 062	03 BSD 024
03 BSC 025	03 BSD 044	03 BSD 064	03 BSD 026
03 BSC 027	03 BSD 048	03 BSD 068	03 BSD 028
03 BSC 029	03 BSD 052	03 BSD 072	03 BSD 032
03 BSC 035	03 BSD 058	03 BSD 078	03 BSD 038

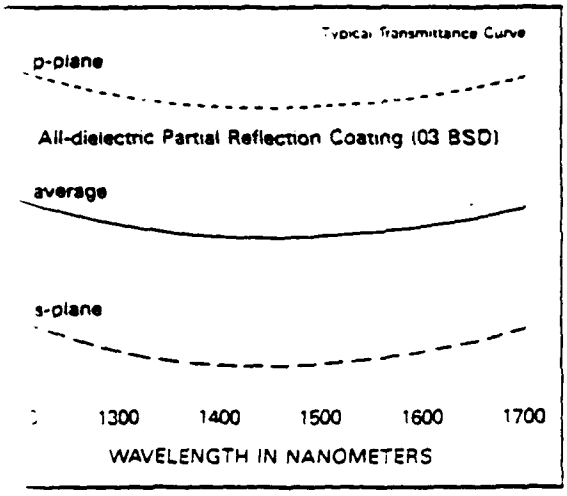

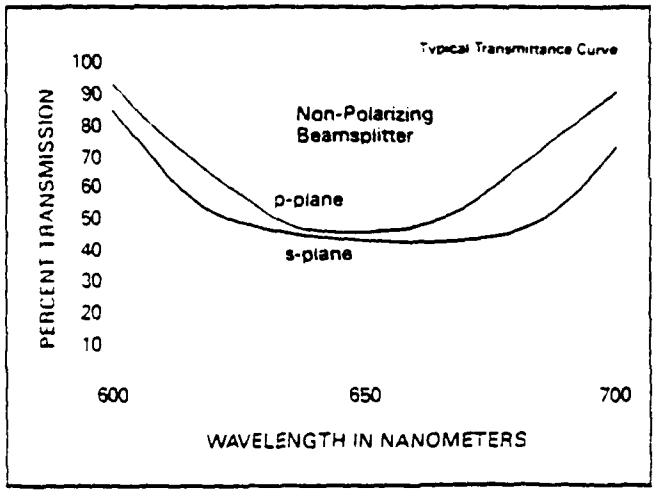

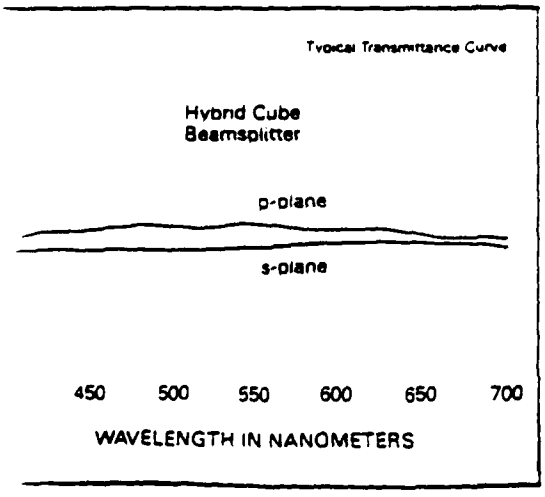

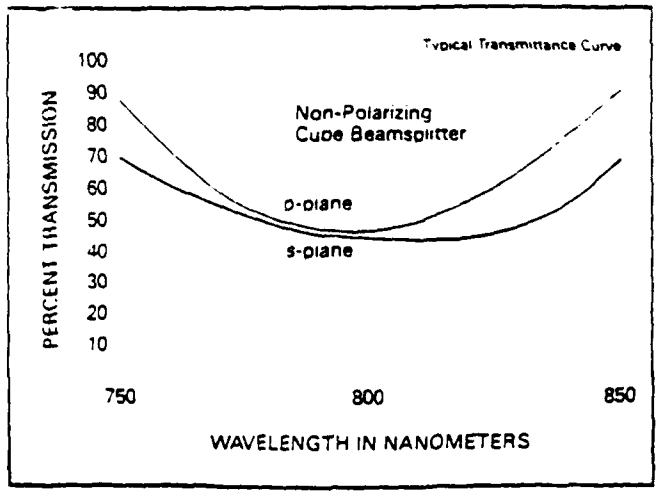

Bibliography

[1] James Burr. Personal conversations throughout 1987–1988. SETI Group, Stanford, California 94043.

[2] Bradley D. Clymer. *Optical Clock Distribution for VLSI.* PhD thesis, Stanford University, Stanford, California, September 1987.

[3] Floyd M. Gardner. *Phaselock Techniques.* John Wiley & Sons, 2nd edition, 1979.

[4] Lance A. Glasser and Daniel W. Dobberpuhl. *The Design and Analysis of VLSI Circuits. VLSI Systems Series*, Addison-Wesley Publishing Company, Inc., December 1985.

[5] G. E. Jellison and F. A. Mondine. Oak Ridge National Laboratory. February 1982.

[6] Carver Mead and Lynn Conway. *Introduction to VLSI Systems.* Addison-Wesley Publishing Company, Inc., Menlo Park, California, 1980.

[7] Masakazu Shoji. Elimination of Process-Dependent Clock Skew in CMOS VLSI. *IEEE Journal of Solid-State Circuits*, SC-21(5):875–880, October 1986.

[8] Ben G. Streetman. *Solid State Electronic Devices.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey 07632, 2nd edition, 1980.

[9] S. M. Sze. *Physics of Semiconductor Devices.* John Wiley & Sons, 2nd edition, 1981.

[10] Neil H. E. Weste and Kamran Eshraghian. *Principles of CMOS VLSI Design.* Addison-Wesley Publishing Company, Inc., Menlo Park, California, 1985.

[11] Bruce A. Wooley. EE313 class notes. Winter Quarter, 1987-88. Stanford University. Stanford, California 94043.